IAMPerformance

Quantum Computing Performance Intelligence

Issue 002 April 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, grounded in information theory -- quantifying how physical systems actualize their computational potential.

NEW IN ISSUE 002	**Position Gauge** -- floor-to-FT visual on every platform card | **Quantinuum Helios** -- full card debut | **Pending cards** -- Microsoft Majorana 1 and Google Neutral Atom | **Section 2: Architecture and Engineering Diagnostics** -- substrate substitution, participation ratio analysis, trajectory acceleration, temperature optimization, combined lever analysis | **THE SUBSTRATE INVERSION** -- a new physical limit derived from first principles: the point at which improving T1 no longer reduces gate error. IBM Nighthawk confirmed past the transition. Google Willow approaching ~2028. Ion trap and neutral atom architectures immune by design. See pp. 28-30. | **Is this the classical Dennard transition in quantum computing?**

14	3	1	5	13
Platforms Tracked	At or Below QEC Threshold	Fault Tolerance Crossed	Active Wall Conditions	Dated Predictions

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

IAMPerformance is an independent research initiative. No institutional affiliation. No investor agenda. No stake in any platform tracked here. The sole mission is to continue refining and optimizing QAPE -- building it into a tool the platforms themselves can use to accelerate the journey toward fault-tolerant quantum computing.

Disclaimer: Not investment advice. For analytical reference only -- not engineering advice. Platform names identify published hardware only. All gate error data from primary publications (cited). Predictions are forward-looking and falsifiable. https://iamperformance.net

IAM ARCHITECTURE CLASSES

Eight architecture classes tracked. The parameter n governs how strongly operating temperature affects gate error rate. Platforms within the same class share the same n -- it is a property of the gate mechanism, not the manufacturer. Pending and pre-commercial classes have n assigned from first principles; departure from assigned n upon first data publication is itself a diagnostic.

Class	Architecture	n	Dominant Noise Channel	Cooling	Platforms
SC CZ	Superconducting, CZ gates	1.449	Flux noise sensitivity	HIGH	Willow, Heron R2, Nighthawk, Cepheus-1 36Q/108Q
SC ECR	Superconducting, iSWAP/ECR gates	1.189	Gate-level decoherence	MODERATE	Ankaa-3 84Q, IBM Eagle
Ion A	Ion trap, laser motional gates	0.612	Electrode field noise	LOW	Helios 98Q, H1-1, H2-1 56Q, IonQ Forte
Ion B	Ion trap, EQC electronic state	1.200	Electronic state noise	MODERATE	IonQ EQC prototype
Neutral	Neutral atom, Rydberg blockade	0.351	Rydberg interaction fidelity	MINIMAL	QuEra Gemini 260Q
NV	NV center, diamond host	2.291	Spin-phonon coupling (diamond)	VERY HIGH	No commercial system yet (highest n)
Topo.	Topological SC (Majorana)	1.189	Gate error unpublished	MODERATE	Microsoft Majorana 1 (pending)
Class G	Photonic, linear optics	0.000	Fabrication loss (photon loss)	NONE	PsiQuantum (not yet published)

THE RACE -- GLOBAL RANKING

Ranked by IAMPerformance A score. Lower is better -- longer bar means better performance. Active platforms ranked. Pending platforms shown below the ranking line.

-- -- QEC threshold (10^-3) -- -- Fault-tolerance target (10^-4)

COMING SOON -- AWAITING FIRST GATE ERROR PUBLICATION

-- Microsoft Majorana 1

COMING SOON | Topological SC (SC ECR class assigned pending gate data)
Majorana 1 | Nature Feb 2025 doi:10.1038/s41586-025-08472-9 | 8 topological qubits | Jul 2025: Z-loop readout error ~0.5%, X-loop ~16% (parity, not gate)
IAM-2026-P006: if A > 0.002 upon first two-qubit gate publication, topological protection is not outperforming SC class. No braiding demonstrated as of Apr 2026.

GATE ERROR PENDING

-- Google Neutral Atom

COMING SOON | Neutral Atom / Rydberg (n = 0.351 assigned)
Google Quantum AI | Announced March 24 2026 | Note: Microsoft + Atom Computing commercial neutral atom system also announced 2025
Architecture class: Neutral Atom. Color code threshold applies (A < 5.0e-3). Two commercial-pathway neutral atom programs now active simultaneously. Framework applied upon first primary-source gate error publication.

GATE ERROR PENDING

PLATFORM DATA -- PRIMARY PUBLISHED SOURCES

All p(2Q) values from primary publications as of April 2026. A score derived directly: A = -ln(1 - p). You can verify any entry: take the published p(2Q), compute -ln(1-p), compare. Sources listed are the specific papers or announcements from which gate error rates are taken. Where multiple chip configurations were characterized, the value used is noted in the platform card.

Platform	Year	p(2Q)	A Score	Primary Source
IonQ EQC prototype	2025	8.40e-5	8.400e-5	arXiv:2510.17286, Hughes et al. (Oxford Ionics/IonQ), Oct 2025

Platform	Year	p(2Q)	A Score	Primary Source
Quantinuum Helios 98Q	2025	7.90e-4	7.903e-4	arXiv:2511.05465, Ransford et al. (Quantinuum), Nov 2025
Quantinuum H1-1	2024	8.60e-4	8.604e-4	Quantinuum blog, Apr 16 2024: 99.914(3)% fidelity
Google Willow 105Q	2024	1.50e-3	1.501e-3	Google Quantum AI, Nature 638:920 (Dec 2024)
Quantinuum H2-1 56Q	2024	1.84e-3	1.842e-3	Quantinuum quantinuum.com: 99.816(5)% fidelity
IBM Heron R2 156Q	2024	2.00e-3	2.002e-3	IBM QDC 2024; AbuGhanem (2024) doi:10.1007/s11227-025-07047-7
IBM Nighthawk 120Q	2026	2.15e-3	2.154e-3	IBM Quantum Platform announcement Jan 5 2026
IonQ Forte 36Q	2024	4.00e-3	4.008e-3	IonQ Forte Enterprise spec ionq.com/quantum-systems/forte-enterprise
Rigetti Cepheus-1 36Q	2025	4.00e-3	4.012e-3	Rigetti Q4 2025 earnings Mar 4 2026: 99.6% median CZ fidelity (36Q system)
QuEra Gemini 260Q	2024	8.00e-3	8.032e-3	QuEra quera.com/gemini: >99.2% two-qubit fidelity
Rigetti Cepheus-1 108Q	2025	1.00e-2	1.005e-2	Rigetti Q4 2025 earnings Mar 4 2026: 99.0% median CZ fidelity (108Q system)
Rigetti Ankaa-3 84Q	2024	1.00e-2	1.005e-2	Rigetti GlobeNewswire Dec 23 2024: 99.0% median iSWAP fidelity
Microsoft Majorana 1	2025	PENDING	PENDING	Aghaee et al., Nature 2025. doi:10.1038/s41586-025-08472-9
Google Neutral Atom (2026+)	--	PENDING	PENDING	Google Quantum AI announcement March 24 2026

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

SECTION 1

CURRENT STATE ANALYSIS

Full IAMPerformance analysis for each active platform in rank order. Every value derived from published gate error rates. Every prediction dated, numbered, and specific.

#1 Oxford Ionics EQC prototype

FAULT TOLERANT

Ion Trap Class B | Ca+ electronic state | arXiv:2510.17286, Hughes et al. (Oxford Ionics/IonQ), Oct 2025

HISTORY

Date	System	Result
Jun 2024	**Oxford Ionics Ca+ 2Q prototype**	p = 3.00e-4 — first published demonstration of the Ca+ electronic-state qubit gate. Two qubits. A proof of concept that the EQC mechanism worked.
Mar 2025	**Oxford Ionics single-qubit record**	p(1Q) = 1.50e-7 — single-qubit world record at the time of publication. Confirmed the architecture could achieve extraordinary fidelity before multi-qubit scaling.
Oct 2025	**IonQ EQC prototype — world #1**	p = 8.40e-5 — two-qubit gate error rate that crossed the fault-tolerance target. No other commercial-pathway system had done this. The paper is arXiv:2510.17286.

PREVIOUS GENERATION vs NOW

THEN Oxford Ionics Ca+ 2Q (Jun 2024) p = 3.0000e-04 **-> NOW** Oxford Ionics EQC prototype p = 8.4000e-05 **(72% improvement)**

IAM read: Rank #1 globally at A=8.4e-5. Ca+ EQC has crossed the FT target — the only commercial-pathway system to do so. 2.80x above Ca+ floor at 0.63-yr halving: wall projected early 2027.

POSITION GAUGE -- Floor | Error Correction | Fault Tolerance

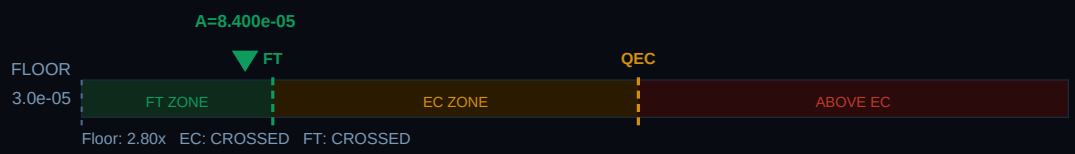

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Status / Source
Published gate error p(2Q)	8.400e-05	**PUBLISHED**
IAMPerformance A score	8.4000e-05	DERIVED: framework metric

Metric	Value	Status / Source
Architecture class	Ion Trap Class B	DERIVED: class assignment
Architecture parameter n	1.200	DERIVED: IAMPerformance
Operating temperature	1 mK	**PUBLISHED**
Improvement rate	0.1 yr (halving)	DERIVED: trajectory
Error correction status	CROSSED	**EC VIABLE**
Fault-tolerance status	CROSSED	**FT VIABLE**
Material substrate	Ca+ electronic state	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 = 8.4e-5, the EQC prototype is the only platform in this dataset to have crossed the fault-tolerance target (A < 1.0e-4). It ranks first globally by a factor of nine over the second-place Helios. The distance to the Ca+ material floor is 2.80x — meaningful headroom at a 0.63-year halving rate that was derived from the Oxford Ca+ 2Q to IonQ EQC improvement trajectory, not assumed. At this rate, the EQC class approaches its Ca+ floor by early 2027. That is approximately twelve months from now. The years-to-FT metric reads CROSSED — this platform can today serve as the foundation for fault-tolerant circuit design. That distinction is shared by no other commercial-pathway system as of April 2026.

COMMENTARY

The IonQ EQC prototype holds the lowest published two-qubit gate error rate of any quantum processor as of April 2026. Not by a little. The paper is arXiv:2510.17286 by Hughes et al. (Oxford Ionics / IonQ, October 2025), and the specific number it gives is 8.4(7)x10^-5 -- that parenthetical seven is the one-sigma uncertainty on the last digit. IonQ's press release described this as exceeding 99.99% fidelity. Both statements are correct. The paper is more precise, and precision is what the IAMPerformance framework runs on. To understand why this result matters so much, you need to understand where IonQ came from. For four years -- from Harmony in 2020 through Aria in 2022 to Forte in 2024 -- IonQ published the same gate error rate: 0.40%. Three systems, four years, identical metric. That is a stall, and the IAMPerformance architecture parameter tells you exactly why it happened. The Yb+ laser motional class has a floor at A = 3.0e-4, and the Harmony-to-Forte line simply reached it and stopped. More engineering did not move the number because the number was already constrained by the underlying physics of ytterbium ions driven by UV lasers. The EQC architecture breaks from this entirely. It uses calcium-43 ions controlled via electronic state transitions rather than the laser-driven motional modes used in the Yb+ class. This is not an incremental improvement on the Forte architecture -- it is a categorically different noise mechanism, Ion Trap Class B versus Class A, with its own architecture parameter (n = 1.200), its own material floor (A = 3.0e-5), and its own improvement trajectory. At A = 8.4e-5, the EQC prototype is 2.80x above the Ca+ floor at a 0.63-year halving rate derived from the Oxford Ca+ 2Q to EQC prototype trajectory. At that rate, the Ca+ material floor is approached by early 2027. The next twelve months will show whether the EQC class can sustain that pace as it scales. Strategic context as of April 2026: IonQ has announced its intention to acquire Oxford Ionics. If completed, the world's highest-performing two-qubit gate platform -- EQC at A = 8.4e-5, fault tolerance already crossed -- moves under IonQ ownership. IonQ's commercial Forte platform sits at rank #8 at A = 4.0e-3, while EQC sits at rank #1 at A = 8.4e-5 -- a factor of 48x better. The acquisition would give IonQ simultaneous ownership of both the best and the most commercially accessible ion trap platforms in the dataset, covering Ion B (EQC) and Ion A (Forte/Aria) architecture classes simultaneously.

TRAJECTORY

Year	A Score	Milestone
2026	8.4000e-05	
2027	3.0000e-05	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
0.25 <- FT	1.5915e-05	0.19x	YES	YES
0.33 <- FT	2.2207e-05	0.26x	YES	YES
0.5 <- FT	3.6563e-05	0.44x	YES	YES
0.67 <- FT	5.1948e-05	0.62x	YES	YES
0.75 <- FT	5.9478e-05	0.71x	YES	YES

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
1.0 <- now	8.4000e-05	--	YES	YES
1.5 <- QEC	1.3664e-04	1.63x	YES	--
2.0 <- QEC	1.9298e-04	2.30x	YES	--
3.0 <- QEC	3.1392e-04	3.74x	YES	--

COOLING EFFECT

HIGH	**Ion Trap Class B n=1.200 1 mK** High temperature sensitivity n=1.200. Halving temp improves error ~2.30x. Cooling has above-average return.

MATERIAL WALL

MEDIUM	**Ca+ electronic state material floor 2.80x above floor A-floor=3.0e-05** Material A-floor is an **IAMPerformance-derived value** -- not published by the manufacturer. The Ca+ EQC floor at A = 3.0e-5 is set by the fundamental noise limit of electronic state transitions in calcium-43 ions controlled by precision microwave electronics. Unlike the laser-driven motional architectures used by Quantinuum and IonQ Forte, EQC has no metal-oxide junction, no laser frequency noise coupling, and no motional heating from electrode field fluctuations as primary error sources. The noise floor here is dominated by electronic phase noise in the microwave control chain and electrode field stability at the semiconductor chip level. At 2.80x above the floor, EQC has meaningful headroom remaining. The improvement levers are electronic control precision, chip fabrication uniformity, and inter-ion crosstalk suppression -- engineering problems, not physics limits. A material change is not required to reach the Ca+ floor. **IAMPerformance predicts:** IAM-2026-P005: EQC approaches Ca+ floor (p~3x10^-5) within 12 months at current rate.

QUANTUM DENNARD TRANSITION

QDT	**RF Control Ceiling** -- **DISTANT (~2035+)** EQC eliminates laser phase noise entirely via RF/microwave control. RF electronics approach shot-noise limit at ~10^-7 single-qubit. Two-qubit at 84 ppm is well above this -- control is not yet limiting. Most engineering runway of any architecture class. The wall is far away.

QUBIT COUNT SCALING LAW

SCL	**FREE** -- alpha ~ 0.00. RF/microwave control has no motional coupling. Most engineering runway of any tracked architecture.

#2 Quantinuum Helios 98Q

QEC NOW

Ion Trap Class A | Ba+ ion (laser) | arXiv:2511.05465, Ransford et al. (Quantinuum), Nov 2025

HISTORY

Date	System	Result
2020	**Quantinuum H1**	p = 3.00e-3 — the Ba+ program baseline. Ion trap, laser motional gates, QCCD architecture. This is where the six-year trajectory begins.
Apr 2024	**Quantinuum H1-1**	p = 8.60e-4 — the commercial QEC threshold crossing. First ion trap system to cross the 1.0e-3 surface code threshold in commercial deployment.
Nov 2025	**Quantinuum Helios 98Q**	p = 7.90e-4 — 74% improvement from H1, 8% improvement from H1-1. 98 qubits. QEC threshold solidly crossed. First full card in IAMPerformance Issue 002.

PREVIOUS GENERATION vs NOW

THEN Quantinuum H1-1 (Apr 2024) p = 8.6000e-04 **-> NOW** Quantinuum Helios 98Q p = 7.9000e-04 **(8% improvement)**

IAM read: Rank #2. QEC crossed. Ba+ floor at 7.9x headroom — architecture not exhausted. The Apollo question: stay on Ba+ or leap to Ca+ EQC?

POSITION GAUGE -- Floor | Error Correction | Fault Tolerance

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Status / Source
Published gate error p(2Q)	7.900e-04	**PUBLISHED**
IAMPerformance A score	7.9030e-04	DERIVED: framework metric
Architecture class	Ion Trap Class A	DERIVED: class assignment
Architecture parameter n	0.612	DERIVED: IAMPerformance
Operating temperature	1 mK	**PUBLISHED**
Improvement rate	2.2 yr (halving)	DERIVED: trajectory
Error correction status	CROSSED	**EC VIABLE**
Fault-tolerance status	7.90x above target	~6.6 yr (2033)
Material substrate	Ba+ ion (laser)	DERIVED: class assignment
Material A-floor	1.0e-04	DERIVED: not published by mfr
Wall concern	7.90x above floor	LOW

GATE ERROR DECOMPOSITION p(2Q) = p_laser + p_motional + p_ctrl (Ion A)

Component	Value (ppm)	% of total	Interpretation
p_laser (phase noise)	620 ppm	78%	Laser quality limited -- improves with laser precision investment
p_motional (heating, cooled)	70 ppm	9%	MHI managed by sympathetic cooling (uncooled: 701 ppm)
p_ctrl (architecture floor)	100 ppm	13%	Irreducible -- Ion A minimum 100 ppm
FIXABLE total (p_laser + p_motional)	690 ppm	87%	Engineering-accessible improvement budget

IAM-ION-P001 (April 3, 2026): CM mode MHI already active without sympathetic cooling (uncooled p_motional = 701 ppm). Helios Yb+ sympathetic cooling suppresses ~10x. Falsifiable: remove sympathetic cooling from a future system and observe regression. Source: arXiv:2206.11888. Patent pending 64/012,720.

METRICS ANALYSIS

Helios sits at A = 7.9e-4, solidly below the QEC threshold of 1.0e-3, ranking second globally. The fault-tolerance target at A < 1.0e-4 has not been crossed — Helios is 7.9x above it. At the 2.2-year Ba+ empirical halving rate, the FT target is approximately 6.6 years away, projecting to around 2032. The 74% improvement from H1 to Helios over six years of Ba+ development is the most consistent improvement record in the ion trap dataset — no regressions, every generation improved. The architecture parameter n = 0.612 tells you directly that cooling is not the accelerator here. Halving the operating temperature from 1 mK to 0.5 mK improves gate error by approximately 1.53x — useful, but the real leverage is in laser pulse optimization and motional mode engineering within the QCCD architecture.

COMMENTARY

Helios arrives in Issue 002 as a full card for the first time -- it was listed as 'emerging' in our Issue 001 baseline, and now the paper has published and the numbers are in. The published two-qubit gate infidelity is 7.9(2)x10^-4 (arXiv:2511.05465, Ransford et al., Quantinuum, November 2025). That places Helios solidly below the QEC threshold of 1.0e-3 and 7.9x above the Ba+ material floor of 1.0e-4. It does not reach the fault-tolerance target. This is still an exceptional result, and reading the paper carefully changes the story that was circulating when it launched. The trajectory that produced Helios is worth understanding. Quantinuum drove from H1 at 3.0e-3 to Helios at 7.9e-4 in six years of continuous Ba+ development -- a 74% total improvement, crossing the commercial QEC threshold with H1-1 along the way, with no regressions and no stalls. The 2.2-year halving rate is the empirical Ba+ program rate derived from that history. It is not a projection from single data points. What the architecture parameter n = 0.612 tells you is that cooling is not Quantinuum's friend here. Halving the operating temperature from 1 mK to 0.5 mK improves gate error by approximately 1.53x -- useful, but not transformative. The leverage is in laser pulse optimization and motional mode engineering within the QCCD architecture, not in the refrigerator. Helios has 7.9x of headroom remaining on the Ba+ floor. The question IAMPerformance is watching for Apollo is whether Quantinuum uses that headroom or leaps to Ca+ EQC now that IonQ has demonstrated what that class can do.

TRAJECTORY

Year	A Score	Milestone
2026	7.9030e-04	
2027	5.7670e-04	
2028	4.2090e-04	
2029	3.0710e-04	
2030	2.2410e-04	
2031	1.6350e-04	
2032	1.1930e-04	
2033	1.0000e-04	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
0.034 <- FT	9.9783e-05	0.13x	YES	YES
0.25 <- QEC	3.3832e-04	0.43x	YES	--
0.33 <- QEC	4.0098e-04	0.51x	YES	--
0.5 <- QEC	5.1708e-04	0.65x	YES	--
0.67 <- QEC	6.1851e-04	0.78x	YES	--
0.75 <- QEC	6.6272e-04	0.84x	YES	--
1.0 <- now	7.9030e-04	--	YES	--
1.5	1.0129e-03	1.28x	--	--
2.0	1.2079e-03	1.53x	--	--
3.0	1.5481e-03	1.96x	--	--

COOLING EFFECT

LOW	**Ion Trap Class A n=0.612 1 mK** Ion trap electrode noise at n=0.612. Halving temp improves error ~1.53x. Gate design is primary lever.

MATERIAL WALL

LOW	**Ba+ ion (laser) material floor 7.90x above floor A-floor=1.0e-04** Material A-floor is an **IAMPerformance-derived value** -- not published by the manufacturer. The Ba+ ion trap floor at A = 1.0e-4 is set by residual motional decoherence of barium-137 hyperfine qubits in the QCCD architecture. Barium's optical transitions at accessible visible wavelengths (493 nm and 650 nm) allow substantially more precise laser control than ytterbium's UV transitions at 369 nm -- this is the core reason Quantinuum moved 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 from here is continued laser pulse optimization and motional mode engineering within the QCCD architecture. No material change is required to close the gap to the Ba+ floor. **IAMPerformance predicts:** IAM-2026-P003: Apollo achieves A < 1.0x10^-4 on Ba+ or switches to Ca+ EQC. | IAM-2026-P011: Motional Heating Inversion (MHI) -- Ion A improvement rate will slow as motional heating from laser power increases becomes comparable to phase noise reduction. DERIVED April 3, 2026: Ba+ CM heating rate 29+/-4 quanta/s (arXiv:2206.11888). Without sympathetic cooling, MHI already active on CM mode. With Helios Yb+ sympathetic cooling (~10x suppression), effective transition ~2033-2037. Falsifiable: remove sympathetic cooling from a future system and observe regression. | IAM-2026-P013: Quantinuum Sol (192Q, 2027): A ~ 5.8x10^-4. QCCD scales for free -- no TLS ceiling, no Substrate Inversion. Sol will NOT regress from Helios.

QUANTUM DENNARD TRANSITION

QDT	**Motional Heating Inversion** -- **MONITORED** Ba+ CM heating rate: 29 ± 4 quanta/s (arXiv:2206.11888) -- CM mode MHI **already active** without sympathetic cooling. Helios Yb+ sympathetic cooling pushes effective transition to ~2033-2037. Escape route implemented. Sol (2027) and Apollo (2029) must maintain it. IAM-ION-P001: gate error will regress if sympathetic cooling is removed from a future system.

QUBIT COUNT SCALING LAW

SCL	**QCCD log(N) -- FREE** -- alpha ~ 0.02. Ion A scales sub-linearly. No substrate ceiling. Sol (192Q, 2027) predicted <0.3% degradation from scaling alone.

#3 Quantinuum H1-1 QEC NOW

Ion Trap Class A | Ba+ ion (laser) | Quantinuum blog, Apr 16 2024: 99.914(3)% fidelity

HISTORY

Date	System	Result
2020	**Quantinuum H1**	p = 3.00e-3 — Ba+ program baseline. Nature 592:209 (Pino et al. 2021).
Apr 2024	**Quantinuum H1-1**	p = 8.60e-4 — 71% improvement from H1. First commercial crossing of the QEC threshold by any ion trap system. The milestone that established Quantinuum as the performance leader in the commercial ion trap class.

PREVIOUS GENERATION vs NOW

THEN Quantinuum H1 (2020) p = 3.0000e-03 **-> NOW** Quantinuum H1-1 p = 8.6000e-04 **(71% improvement)**

IAM read: Rank #3. QEC crossed April 2024 — still the historical milestone for commercial ion trap. 8.6x above Ba+ floor with 2.2-yr halving: wall ~2033. The platform that proved it was possible.

POSITION GAUGE -- Floor | Error Correction | Fault Tolerance

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Status / Source
Published gate error p(2Q)	8.600e-04	**PUBLISHED**
IAMPerformance A score	8.6040e-04	DERIVED: framework metric
Architecture class	Ion Trap Class A	DERIVED: class assignment
Architecture parameter n	0.612	DERIVED: IAMPerformance
Operating temperature	1 mK	**PUBLISHED**
Improvement rate	6.5 yr (halving)	DERIVED: trajectory
Error correction status	CROSSED	**EC VIABLE**
Fault-tolerance status	8.60x above target	~20.2 yr (2047)
Material substrate	Ba+ ion (laser)	DERIVED: class assignment
Material A-floor	1.0e-04	DERIVED: not published by mfr
Wall concern	8.60x above floor	LOW

GATE ERROR DECOMPOSITION p(2Q) = p_laser + p_motional + p_ctrl (Ion A)

Component	Value (ppm)	% of total	Interpretation
p_laser (phase noise)	690 ppm	80%	Laser quality limited -- improves with laser precision investment
p_motional (heating, cooled)	70 ppm	8%	MHI managed by sympathetic cooling (uncooled: 701 ppm)
p_ctrl (architecture floor)	100 ppm	12%	Irreducible -- Ion A minimum 100 ppm
FIXABLE total (p_laser + p_motional)	760 ppm	88%	Engineering-accessible improvement budget

IAM-ION-P001 (April 3, 2026): CM mode MHI already active without sympathetic cooling (uncooled p_motional = 701 ppm). Helios Yb+ sympathetic cooling suppresses ~10x. Falsifiable: remove sympathetic cooling from a future system and observe regression. Source: arXiv:2206.11888. Patent pending 64/012,720.

METRICS ANALYSIS

H1-1 established the commercial QEC threshold crossing in April 2024 at A = 8.6e-4. It ranks third globally and sits clearly below the QEC threshold. Distance to FT is 8.6x. The 2.2-year Ba+ halving rate reflects the observed 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 6.8 years, projecting to around 2033. The architecture parameter n = 0.612 applies here exactly as it does for Helios: cooling returns modest improvement, and gate design is the primary lever for any further progress.

COMMENTARY

H1-1 was the first commercial crossing of the QEC threshold by any ion trap system. That milestone happened in April 2024 and it still holds its historical significance -- the IonQ EQC prototype later crossed the fault-tolerance target, and Helios sits right alongside H1-1, but H1-1 was where the commercial error correction threshold went from theory to hardware. That is worth remembering every time you look at this card. The system currently sits at A = 8.6e-4, 8.6x above the Ba+ material floor, with a 2.2-year halving rate that puts the wall around 2033. The architecture parameter n = 0.612 applies here exactly as it does for Helios: cooling from 1 mK to 0.5 mK produces approximately 1.53x improvement in gate error rate -- useful, but not the primary lever. Gate design and laser control are where H1-1's next improvements come from, not the refrigerator. H1-1 remains commercially deployed. It continues to define the accessible QEC-viable baseline for the algorithmic research community -- the platform researchers reach for when they want to run error-corrected circuits on hardware that is actually available today.

TRAJECTORY

Year	A Score	Milestone
2026	8.6040e-04	
2027	7.7330e-04	
2028	6.9510e-04	
2029	6.2480e-04	
2030	5.6160e-04	
2031	5.0480e-04	
2032	4.5370e-04	
2033	4.0780e-04	
2034	3.6660e-04	
2035	3.2950e-04	
2036	2.9620e-04	
2037	2.6620e-04	
2038	2.3930e-04	
2039	2.1510e-04	
2040	1.9330e-04	
2041	1.7380e-04	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
0.029 <- FT	9.8557e-05	0.11x	YES	YES
0.25 <- QEC	3.6833e-04	0.43x	YES	--
0.33 <- QEC	4.3655e-04	0.51x	YES	--
0.5 <- QEC	5.6295e-04	0.65x	YES	--
0.67 <- QEC	6.7338e-04	0.78x	YES	--
0.75 <- QEC	7.2150e-04	0.84x	YES	--
1.0 <- now	8.6040e-04	--	YES	--
1.5	1.1027e-03	1.28x	--	--
2.0	1.3150e-03	1.53x	--	--
3.0	1.6854e-03	1.96x	--	--

COOLING EFFECT

LOW	**Ion Trap Class A n=0.612 1 mK** Ion trap electrode noise at n=0.612. Halving temp improves error ~1.53x. Gate design is primary lever.

MATERIAL WALL

LOW	**Ba+ ion (laser) material floor 8.60x above floor A-floor=1.0e-04** Material A-floor is an **IAMPerformance-derived value** -- not published by the manufacturer. H1-1 shares the Ba+ ion trap architecture and the same floor (A = 1.0e-4) with Helios. The physical limit is identical: residual motional decoherence of Ba+ hyperfine qubits in QCCD geometry. 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, along exactly the same path that produced the H1-to-H1-1-to-Helios trajectory. **IAMPerformance predicts:** IAM-2026-P003: Apollo achieves A < 1.0x10^-4 on Ba+ or switches to Ca+ EQC. | IAM-2026-P011: Motional Heating Inversion (MHI) -- Ion A improvement rate will slow as motional heating from laser power increases becomes comparable to phase noise reduction. DERIVED April 3, 2026: Ba+ CM heating rate 29+/-4 quanta/s (arXiv:2206.11888). Without sympathetic cooling, MHI already active on CM mode. With Helios Yb+ sympathetic cooling (~10x suppression), effective transition ~2033-2037. Falsifiable: remove sympathetic cooling from a future system and observe regression. | IAM-2026-P013: Quantinuum Sol (192Q, 2027): A ~ 5.8x10^-4. QCCD scales for free.

QUANTUM DENNARD TRANSITION

QDT	**Motional Heating Inversion** -- **MONITORED** Ba+ CM heating rate: 29 ± 4 quanta/s (arXiv:2206.11888) -- CM mode MHI **already active** without sympathetic cooling. Helios Yb+ sympathetic cooling pushes effective transition to ~2033-2037. Escape route implemented. Sol (2027) and Apollo (2029) must maintain it. IAM-ION-P001: gate error will regress if sympathetic cooling is removed from a future system.

QUBIT COUNT SCALING LAW

SCL	**QCCD log(N) -- FREE** -- alpha ~ 0.02. Ion A scales sub-linearly. No substrate ceiling. Sol (192Q, 2027) predicted <0.3% degradation from scaling alone.

#4 Google Willow 105Q

HIGH

SC CZ Gates | AI+PR-engineering | Google Quantum AI, Nature 638:920 (Dec 2024)

HISTORY

Date	System	Result
2018	**Google Bristlecone 72Q**	p ~ 3.00e-2 — class-level estimate. The Al/AlOx transmon era begins. No peer-reviewed gate error rate published for Bristlecone.
Sep 2019	**Google Sycamore 53Q**	p = 6.00e-3 — Nature 574:505 (Arute et al.). The quantum supremacy paper. Al/AlOx CZ-gate transmon at its best-case native performance.
2019-2024	**FIVE-YEAR PLATEAU**	Near-zero improvement on standard AlOx floor. The material ceiling had been reached. Engineering effort continued but the gate error barely moved. This is what a material wall looks like in the data.
Dec 2024	**Google Willow 105Q**	p = 1.50e-3 — Nature 638:920. Participation ratio engineering broke the five-year plateau. T2/T1 = 1.309 flagged as anomalous. Two chip configurations characterized.

PREVIOUS GENERATION vs NOW

THEN Google Sycamore 53Q (2019) p = 6.0000e-03 **-> NOW** Google Willow 105Q p = 1.5000e-03 **(75% improvement)**

IAM read: Rank #4. QEC not yet crossed on gate error alone. PR engineering ceiling reached — 1.50x above floor. The next chip needs tantalum. Cooling to 7.5 mK crosses QEC without a new processor.

POSITION GAUGE -- Floor | Error Correction | Fault Tolerance

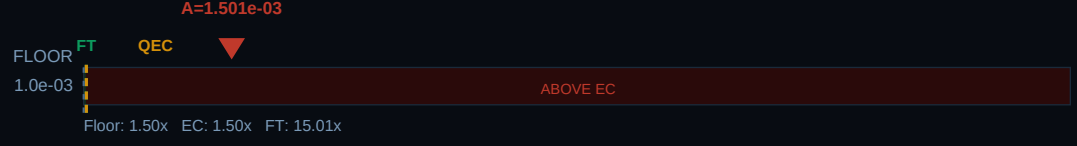

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Status / Source
Published gate error p(2Q)	1.500e-03	**PUBLISHED**
IAMPerformance A score	1.5011e-03	DERIVED: framework metric
Architecture class	SC CZ Gates	DERIVED: class assignment
Architecture parameter n	1.449	DERIVED: IAMPerformance
Operating temperature	15 mK	**PUBLISHED**
Improvement rate	3.3 yr (halving)	DERIVED: trajectory
Error correction status	1.50x above threshold	~1.9 yr (2028)
Fault-tolerance status	15.01x above target	~12.9 yr (2039)
Material substrate	Al+PR-engineering	DERIVED: class assignment
Material A-floor	1.0e-03	DERIVED: not published by mfr
Wall concern	1.50x above floor	STALL
T1 coherence time	68.0 us	**PUBLISHED**
Substrate Inversion T1*	T1* ~ 325 us (T1_free=500 us)	PRE-TRANSITION (T1/T1*=0.21x, T1*=325 us)
T2/T1 ratio	1.309	**ANOMALOUS (!) -- T2 > T1**

GATE ERROR DECOMPOSITION p(2Q) = p_coh + p_mat + p_ctrl

Component	Value (ppm)	% of total	Interpretation
p_coh (coherence)	907 ppm	60%	T1/T2 limited -- improves with coherence time
p_mat (substrate TLS)	56 ppm	4%	TLS defects -- T1_free ceiling = 500 us
p_ctrl (architecture floor)	537 ppm	36%	Irreducible -- architecture minimum 550 ppm
FIXABLE total (p_coh + p_mat)	963 ppm	64%	Engineering-accessible improvement budget

METRICS ANALYSIS

Willow sits at A = 1.501e-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 CZ-gate SC platforms. The 3.3-year halving rate projects QEC crossing around 2028 at current pace — assuming no material ceiling intervenes. The thermal prediction provides a faster path: at 7.5 mK, IAMPerformance predicts p(2Q) = 5.50e-4, crossing QEC without a new processor generation. That is a specific, falsifiable claim with a number attached to it.

COMMENTARY

The Sycamore-to-Willow gap is five years of near-zero improvement. That gap is important to understand before you look at the Willow number, because it tells you exactly what was happening: the standard AlOx amorphous oxide floor was the ceiling, and Sycamore had reached it in 2019. Five years of engineering effort -- new qubit geometries, better calibration, improved fabrication -- moved the gate error barely at all, because the number was constrained by the underlying material, not by the engineering. Willow broke through that ceiling via participation ratio engineering -- a geometric technique that reduces the fraction of the qubit's electric field energy stored in the lossy oxide layer, without changing the oxide material itself. T1 improved from approximately 20 us on Sycamore to approximately 100 us on Willow (Chip 2 RCS: 98 us +/- 32 us). A factor of five. Willow is not on a new material. It is on a more cleverly engineered version of the same material, which is why the A-floor assigned here is 1.0e-3 rather than the 3.0e-3 naive AlOx floor. Willow at A = 1.501e-3 is 1.50x above that engineered floor. A note on gate types that matters for this card: Willow was characterized on two chip configurations. Chip 1 (used for QEC demonstrations) uses CZ gates and reports 0.33% +/- 0.18% two-qubit error. Chip 2 (used for random circuit sampling) uses an iSWAP-like gate and reports 0.14% +/- 0.05% -- the best published two-qubit gate error for Willow. The p = 1.5e-3 used here is the median from Chip 2's iSWAP-like benchmark. 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 IAMPerformance finds 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 = 5.50e-4 -- below the QEC threshold -- without building a new chip. That is a specific, falsifiable prediction with a number attached to it. 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 -- and 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 beyond the floor that participation ratio engineering has now reached.

TRAJECTORY

Year	A Score	Milestone
2026	1.5011e-03	
2027	1.2167e-03	
2028	1.0000e-03	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
2.29 <- FT	9.8550e-05	0.07x	YES	YES
3.75 <- QEC	2.0138e-04	0.13x	YES	--
4.95 <- QEC	3.0112e-04	0.20x	YES	--
7.5 <- QEC	5.4982e-04	0.37x	YES	--
10.05 <- QEC	8.4022e-04	0.56x	YES	--
11.22 <- QEC	9.8558e-04	0.66x	YES	--
11.25 <- QEC	9.8940e-04	0.66x	YES	--
15.0 <- now	1.5011e-03	--	--	--
22.5	2.7013e-03	1.80x	--	--
30.0	4.0983e-03	2.73x	--	--
45.0	7.3749e-03	4.92x	--	--

COOLING EFFECT

HIGH	**SC CZ Gates n=1.449 15 mK** High temperature sensitivity n=1.449. Halving temp improves error ~2.73x. Cooling has above-average return.

MATERIAL WALL

STALL	**Al+PR-engineering material floor 1.50x above floor A-floor=1.0e-03** Material A-floor is an **IAMPerformance-derived value** -- not published by the manufacturer. Willow's material floor at A = 1.0e-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 itself. 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 -- by reducing the fraction of the qubit's electric field stored in the lossy oxide -- but did not eliminate the defects physically. Willow at 1.50x above this engineered floor has essentially reached it. Tantalum junctions have a significantly lower TLS defect density than AlOx: tantalum's crystalline Ta2O5 oxide layer forms a fundamentally cleaner interface than amorphous AlOx. Several academic groups have demonstrated tantalum-based transmons with error rates approaching 1e-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, p(2Q)=5.50x10^-4 -- QEC crossed without a new chip.

QUANTUM DENNARD TRANSITION

QDT	**Substrate Inversion** -- **PRE-TRANSITION** T1 = 68.0 µs · T1* = 325 µs · T1/T1* = 0.21x · T1_free (Al+PR-engineering) = 500 µs Significant T1 improvement runway remains before hitting the Substrate Inversion. Engineering investment in coherence will continue to pay off.

QUBIT COUNT SCALING LAW

SCL	**PRE-T1* (FREE)** -- alpha ~ 0.00. Qubit count scaling essentially free at current substrate state.

#5 Quantinuum H2-1 56Q

HIGH

Ion Trap Class A | Ba+ ion (laser) | Quantinuum quantinuum.com: 99.816(5)% fidelity

HISTORY

Date	System	Result
2020	**Quantinuum H1**	p = 3.00e-3 — Ba+ program baseline, Nature 592:209.
2024	**Quantinuum H2-1 56Q**	p = 1.84e-3 — 38% improvement from H1. 56 qubits, all-to-all connectivity within zones. The commercial workhorse of the H-series line.

PREVIOUS GENERATION vs NOW

THEN Quantinuum H1 (2020) p = 3.0000e-03 **-> NOW** Quantinuum H2-1 56Q p = 1.8400e-03 **(39% improvement)**

IAM read: Rank #5. Commercial workhorse. 18.4x above Ba+ floor -- most headroom of any Quantinuum platform. Roadmap: Helios (deployed) -> Sol (192Q, 2027) -> Apollo (FT, 2029). MHI flag: sympathetic cooling is the countermeasure that sustains this trajectory.

POSITION GAUGE -- Floor | Error Correction | Fault Tolerance

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Status / Source
Published gate error p(2Q)	1.840e-03	**PUBLISHED**
IAMPerformance A score	1.8417e-03	DERIVED: framework metric
Architecture class	Ion Trap Class A	DERIVED: class assignment
Architecture parameter n	0.612	DERIVED: IAMPerformance
Operating temperature	1 mK	**PUBLISHED**
Improvement rate	6.5 yr (halving)	DERIVED: trajectory
Error correction status	1.84x above threshold	~5.7 yr (2032)
Fault-tolerance status	18.42x above target	~27.3 yr (2054)
Material substrate	Ba+ ion (laser)	DERIVED: class assignment
Material A-floor	1.0e-04	DERIVED: not published by mfr
Wall concern	18.42x above floor	LOW

GATE ERROR DECOMPOSITION p(2Q) = p_laser + p_motional + p_ctrl (Ion A)

Component	Value (ppm)	% of total	Interpretation
p_laser (phase noise)	1670 ppm	91%	Laser quality limited -- improves with laser precision investment
p_motional (heating, cooled)	70 ppm	4%	MHI managed by sympathetic cooling (uncooled: 701 ppm)
p_ctrl (architecture floor)	100 ppm	5%	Irreducible -- Ion A minimum 100 ppm
FIXABLE total (p_laser + p_motional)	1740 ppm	95%	Engineering-accessible improvement budget

IAM-ION-P001 (April 3, 2026): CM mode MHI already active without sympathetic cooling (uncooled p_motional = 701 ppm). Helios Yb+ sympathetic cooling suppresses ~10x. Falsifiable: remove sympathetic cooling from a future system and observe regression. Source: arXiv:2206.11888. Patent pending 64/012,720.

METRICS ANALYSIS

H2-1 sits at A = 1.842e-3, above the QEC threshold, ranking fifth globally. Distance to QEC is 1.84x. At 18.4x above the Ba+ floor, H2-1 has the most material headroom of any Quantinuum platform — but the improvement rate reflects that the commercial system is not the performance priority. The slow halving rate is not constrained by physics. H2-1's 56-qubit all-to-all connectivity within zones means its algorithmic value is not fully 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.

COMMENTARY

H2-1 is not the Quantinuum performance leader and is not trying to be -- and that is worth stating clearly up front, because looking at the A score without that context could give the wrong impression. H2-1 is the commercial system: 56 qubits, all-to-all connectivity within zones, running customer workloads today. The 2.2-year Ba+ program halving rate reflects that the engineering priority is on the performance systems (Helios, Sol, Apollo) rather than on pushing H2-1's gate error. That is a business decision, not a physics limit. At 18.4x above the Ba+ material floor, H2-1 has no wall concern on the material side. The commercial system is stable. The architecture parameter n = 0.612 applies here exactly as it does across the Quantinuum ion trap family: ion trap systems operate at room temperature with laser-cooled ions -- cooling investment returns approximately 35% improvement at best (halving T improves gate error by 1.53x), but this is not the engineering lever that matters for H2-1. The investment that matters is on software and systems -- TKET compilation efficiency, H-series connectivity architecture, the hybrid classical-quantum stack. H2-1's 56-qubit all-to-all connectivity also means its algorithmic value is not fully captured by the gate error metric 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. The Ba+ roadmap is now fully defined: Helios (98Q, 2025, now in deployment), Sol (192Q, 2027), and Apollo (fully fault-tolerant, 2029) -- a direct line from H2-1 through each generation on the same ion trap architecture. The IAMPerformance framework places one technical flag on this trajectory: the Motional Heating Inversion. Without sympathetic cooling, the Ba+ CM heating rate of 29 quanta/second already exceeds the laser phase noise floor at current performance. Helios addressed this with Yb+ sympathetic cooling. Sol and Apollo will need to maintain or extend that countermeasure as qubit count scales. The physics is not the obstacle. The architecture already knows the answer. H2-1 is not the story -- it is the baseline the story started from.

TRAJECTORY

Year	A Score	Milestone
2026	1.8417e-03	
2027	1.6554e-03	
2028	1.4880e-03	
2029	1.3375e-03	
2030	1.2022e-03	
2031	1.0806e-03	
2032	9.7130e-04	
2033	8.7300e-04	
2034	7.8470e-04	
2035	7.0540e-04	
2036	6.3400e-04	
2037	5.6990e-04	
2038	5.1220e-04	
2039	4.6040e-04	
2040	4.1390e-04	
2041	3.7200e-04	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
0.008 <- FT	9.5920e-05	0.05x	YES	YES
0.25 <- QEC	7.8842e-04	0.43x	YES	--
0.33 <- QEC	9.3444e-04	0.51x	YES	--
0.365 <- QEC	9.9390e-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	--	--	--
1.5	2.3604e-03	1.28x	--	--

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
2.0	2.8148e-03	1.53x	--	--
3.0	3.6076e-03	1.96x	--	--

COOLING EFFECT

LOW	**Ion Trap Class A n=0.612 1 mK** Ion trap electrode noise at n=0.612. Halving temp improves error ~1.53x. Gate design is primary lever.

MATERIAL WALL

LOW	**Ba+ ion (laser) material floor 18.42x above floor A-floor=1.0e-04** Material A-floor is an **IAMPerformance-derived value** -- not published by the manufacturer. H2-1 shares the Ba+ ion trap floor at A = 1.0e-4 with Helios and H1-1. At 18.4x above the floor, H2-1 has the most material headroom of any Quantinuum platform in this dataset -- the slow improvement rate reflects deployment priorities, not physics limits. The material floor is a long way away. The binding technical question for the Ba+ roadmap is not the material floor but the Motional Heating Inversion: at 29 quanta/second CM heating rate, sympathetic cooling is the countermeasure that keeps this architecture on its improvement trajectory through Sol (2027) and Apollo (2029). **IAMPerformance predicts:** IAM-2026-P003: Apollo achieves A < 1.0x10^-4 on Ba+ or switches to Ca+ EQC. | IAM-2026-P011: Motional Heating Inversion (MHI) -- Ion A improvement rate will slow as motional heating from laser power increases becomes comparable to phase noise reduction. DERIVED April 3, 2026: Ba+ CM heating rate 29+/-4 quanta/s (arXiv:2206.11888). Without sympathetic cooling, MHI already active on CM mode. With Helios Yb+ sympathetic cooling (~10x suppression), effective transition ~2033-2037. Falsifiable: remove sympathetic cooling from a future system and observe regression. | IAM-2026-P013: Quantinuum Sol (192Q, 2027): A ~ 5.8x10^-4. QCCD scales for free.

QUANTUM DENNARD TRANSITION

QDT	**Motional Heating Inversion** -- **MONITORED** Ba+ CM heating rate: 29 ± 4 quanta/s (arXiv:2206.11888) -- CM mode MHI **already active** without sympathetic cooling. Helios Yb+ sympathetic cooling pushes effective transition to ~2033-2037. Escape route implemented. Sol (2027) and Apollo (2029) must maintain it. IAM-ION-P001: gate error will regress if sympathetic cooling is removed from a future system.

QUBIT COUNT SCALING LAW

SCL	**QCCD log(N) -- FREE** -- alpha ~ 0.02. Ion A scales sub-linearly. No substrate ceiling. Sol (192Q, 2027) predicted <0.3% degradation from scaling alone.

#6 IBM Heron R2 156Q

HIGH

SC CZ Gates | Nb+AlOx | IBM QDC 2024; AbuGhanem (2024) doi:10.1007/s11227-025-07047-7

HISTORY

Date	System	Result
2021	**IBM Eagle 127Q**	p = 1.00e-2 — SC ECR gate class. Nb+AlOx substrate. The starting point for the IBM CZ transition.
2023	**IBM Heron R1**	Architecture class change: ECR gates to CZ gates. First IBM processor to move to flux-tunable couplers. The class change shifted the dominant noise mechanism and the architecture parameter.
2024	**IBM Heron R2 156Q**	p = 2.00e-3 — 80% improvement from Eagle over three years. CZ architecture mature. TLS mitigation features introduced. The best IBM result on the Nb+AlOx substrate.

PREVIOUS GENERATION vs NOW

THEN IBM Eagle 127Q (2021) p = 1.0000e-02 **-> NOW** IBM Heron R2 156Q p = 2.0000e-03 **(80% improvement)**

IAM read: Rank #6. 80% improvement from Eagle, but the Nb+AlOx wall is here at 2.00x above floor. TLS mitigation has given what it can give. IBM Starling needs a different path.

POSITION GAUGE -- Floor | Error Correction | Fault Tolerance

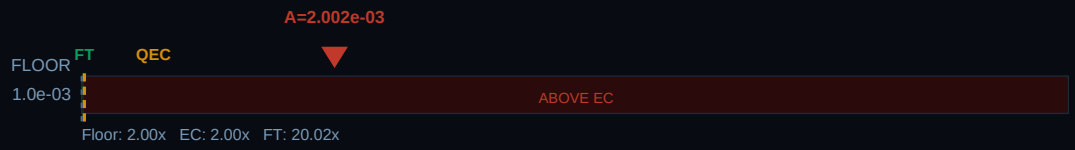

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Status / Source
Published gate error p(2Q)	2.000e-03	**PUBLISHED**
IAMPerformance A score	2.0020e-03	DERIVED: framework metric
Architecture class	SC CZ Gates	DERIVED: class assignment
Architecture parameter n	1.449	DERIVED: IAMPerformance
Operating temperature	15 mK	**PUBLISHED**
Improvement rate	2.0 yr (halving)	DERIVED: trajectory
Error correction status	2.00x above threshold	~2.0 yr (2029)
Fault-tolerance status	20.02x above target	~8.6 yr (2035)
Material substrate	Nb+AlOx	DERIVED: class assignment
Material A-floor	1.0e-03	DERIVED: not published by mfr
Wall concern	2.00x above floor	HIGH
T1 coherence time	300.0 us	**PUBLISHED**
Substrate Inversion T1*	T1* ~ 325 us (T1_free=500 us)	APPROACHING T1* (T1/T1*=0.92x, T1*=325 us)

GATE ERROR DECOMPOSITION p(2Q) = p_coh + p_mat + p_ctrl

Component	Value (ppm)	% of total	Interpretation
p_coh (coherence)	250 ppm	12%	T1/T2 limited -- improves with coherence time
p_mat (substrate TLS)	125 ppm	6%	TLS defects -- T1_free ceiling = 500 us
p_ctrl (architecture floor)	1625 ppm	81%	Irreducible -- architecture minimum 550 ppm
FIXABLE total (p_coh + p_mat)	375 ppm	19%	Engineering-accessible improvement budget

METRICS ANALYSIS

Heron R2 sits at A = 2.002e-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 = 1.0e-2 to 2.0e-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. SCALING LAW CAVEAT: Heron R2 sits at T1/T1* = 0.92x -- approaching the Substrate Inversion. The 2.0-year halving rate from the Heron generation forward is unreliable as a forward predictor if qubit count increases on Nb+AlOx. The Nighthawk result already demonstrates the Substrate Inversion Amplifier engaging: +9% connectivity caused +7.6% regression. The trajectory table assumes the halving rate continues -- it will not if the next IBM generation scales connectivity on the same substrate. Tantalum changes this entirely.

COMMENTARY

The Eagle-to-Heron transition is not just an incremental improvement -- it is an architecture class change, and understanding that change is the key to understanding where IBM stands today. IBM moved from ECR gates to CZ gates, which shifted the dominant noise mechanism and the architecture parameter from n = 1.189 to n = 1.449. That shift means the IBM SC class is now significantly more temperature-sensitive than it was on the Eagle architecture, and the thermal predictions on this card reflect that. The 80% improvement from Eagle to Heron R2 -- from A = 1.0e-2 to 2.0e-3 over three years -- reflects both the class change and four years of engineering investment, including TLS mitigation features introduced in Heron R2. TLS mitigation is a direct attack on the two-level system losses that define the Nb+AlOx oxide floor: it controls the electrostatic environment around junctions to reduce resonant interactions with TLS defects. Despite that effort, Heron R2 sits at 2.0x above the Nb+AlOx material 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. That six-year gap is the acceleration that the substrate has to provide -- or that a material change would need to close. The thermal prediction is the bridge: cooling Heron R2 to 7.5 mK is predicted at p = 7.33e-4, which crosses QEC. That is a cryogenic infrastructure decision, not a fabrication decision, and with n = 1.449 the return on cooling investment for this architecture is real.

TRAJECTORY

Year	A Score	Milestone
2026	2.0020e-03	
2027	1.4156e-03	
2028	1.0010e-03	
2029	1.0000e-03	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
1.877 <- FT	9.8528e-05	0.05x	YES	YES
3.75 <- QEC	2.6858e-04	0.13x	YES	--
4.95 <- QEC	4.0160e-04	0.20x	YES	--
7.5 <- QEC	7.3328e-04	0.37x	YES	--
9.198 <- QEC	9.8560e-04	0.49x	YES	--
10.05	1.1206e-03	0.56x	--	--
11.25	1.3196e-03	0.66x	--	--
15.0 <- now	2.0020e-03	--	--	--
22.5	3.6026e-03	1.80x	--	--
30.0	5.4658e-03	2.73x	--	--
45.0	9.8359e-03	4.92x	--	--

COOLING EFFECT

HIGH

SC CZ Gates n=1.449 15 mK
High temperature sensitivity n=1.449. Halving temp improves error ~2.73x. Cooling has above-average return.

MATERIAL WALL

HIGH

Nb+AlOx material floor 2.00x above floor A-floor=1.0e-03
Material A-floor is an **IAMPerformance-derived value** -- not published by the manufacturer.

IBM's Nb+AlOx floor at A = 1.0e-3 arises from the same TLS defect mechanism as Google's AlOx junction, but with niobium as the superconducting base layer. Niobium's higher critical temperature (9.2 K versus 1.2 K for aluminum) and longer baseline coherence explain why IBM niobium transmons historically outperformed aluminum transmons -- but the AlOx junction barrier is shared between both, and that is where TLS defects live. IBM's TLS mitigation feature in Heron R2 controls the electrostatic environment around junctions to reduce resonant TLS interactions. At 2.0x above the Nb+AlOx floor, Heron R2 has extracted most of what TLS mitigation can provide on this substrate. The Nighthawk regression confirmed this: more connectivity pushed the effective TLS coupling up, not down.
IAMPerformance predicts: IAM-2026-P002: A < 0.002 not achievable on Nb+AlOx without material change. | IAM-2026-P012: Next IBM generation on Nb/AlOx (N > 218Q) will regress further -- A > 2.3x10^-3. Substrate Inversion Amplifier is engaged. Tantalum returns scaling to free.

QUANTUM DENNARD TRANSITION

QDT

Substrate Inversion -- **APPROACHING T1***
T1 = 300.0 µs · T1* = 325 µs · T1/T1* = 0.92x · T1_free (Nb+AlOx) = 500 µs
Within 1-2 chip generations of the Substrate Inversion. The next platform should begin substrate transition planning now. T1 improvement budget is narrowing.

QUBIT COUNT SCALING LAW

SCL

PRE-T1* (FREE) -- alpha ~ 0.00. Qubit count scaling essentially free at current substrate state.

#7 IBM Nighthawk 120Q

HIGH

SC CZ Gates | Nb+AlOx | IBM Quantum Platform announcement Jan 5 2026

HISTORY

Date	System	Result
2021	**IBM Eagle 127Q**	p = 1.00e-2 — the baseline for this IBM trajectory.
2024	**IBM Heron R2 156Q**	p = 2.00e-3 — best IBM result before Nighthawk. 156 qubits, 176 couplers.
Jan 2026	**IBM Nighthawk 120Q**	p = 2.15e-3 — gate error INCREASED from Heron R2. 218 couplers. T1 = 350 us (highest in IBM fleet). More connectivity, better coherence, worse gate error. This is what a material wall looks like when you push past it.

PREVIOUS GENERATION vs NOW

THEN IBM Heron R2 156Q (2024) p = 2.0000e-03 **->** **NOW** IBM Nighthawk 120Q p = 2.1540e-03 **(8% REGRESSION)**

IAM read: Rank #7 — behind Heron R2. Error rate INCREASED. T1 best in fleet. The material ceiling confirmed: more connectivity on Nb+AlOx pushes past the junction quality limit.

POSITION GAUGE -- Floor | Error Correction | Fault Tolerance

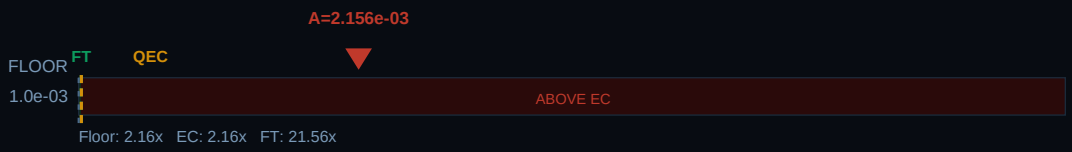

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Status / Source
Published gate error p(2Q)	2.154e-03	**PUBLISHED**
IAMPerformance A score	2.1563e-03	DERIVED: framework metric
Architecture class	SC CZ Gates	DERIVED: class assignment
Architecture parameter n	1.449	DERIVED: IAMPerformance
Operating temperature	15 mK	**PUBLISHED**
Improvement rate	2.0 yr (halving)	DERIVED: trajectory
Error correction status	2.16x above threshold	~2.2 yr (2029)
Fault-tolerance status	21.56x above target	~8.9 yr (2035)
Material substrate	Nb+AlOx	DERIVED: class assignment
Material A-floor	1.0e-03	DERIVED: not published by mfr
Wall concern	2.16x above floor	HIGH
T1 coherence time	350.0 us	**PUBLISHED**
Substrate Inversion T1*	T1* ~ 325 us (T1_free=500 us)	PAST T1* -- IN TRANSITION (T1/T1*=1.08x)

GATE ERROR DECOMPOSITION p(2Q) = p_coh + p_mat + p_ctrl

Component	Value (ppm)	% of total	Interpretation
p_coh (coherence)	214 ppm	10%	T1/T2 limited -- improves with coherence time
p_mat (substrate TLS)	167 ppm	8%	TLS defects -- T1_free ceiling = 500 us
p_ctrl (architecture floor)	1773 ppm	82%	Irreducible -- architecture minimum 550 ppm
FIXABLE total (p_coh + p_mat)	381 ppm	18%	Engineering-accessible improvement budget

METRICS ANALYSIS

Nighthawk sits at A = 2.152e-3, ranking seventh -- behind its predecessor Heron R2. This is the only platform in this dataset where a successor has a higher A score than the system it followed. Distance to QEC is 2.15x. T1 = 350 us 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 topology on the same Nb+AlOx substrate. The material wall is real and Nighthawk is its clearest confirmation. SCALING LAW CAVEAT: The 2.0-year halving rate trajectory on this card is not a reliable forward predictor. Nighthawk is past T1* (T1/T1* = 1.08x). The Substrate Inversion Amplifier is active: the crosstalk cost per added coupler is now substantially larger than it was in the Falcon-to-Eagle era -- a consequence of the TLS defect environment being saturated at every junction. Any next IBM generation that increases connectivity on Nb+AlOx will regress further, not improve. The trajectory assumes the halving rate resumes -- that requires either a substrate change (Tantalum: T1* = 1300 us) or a gate mechanism redesign that bypasses the T1/substrate relationship entirely. IAM-2026-P012: next IBM generation on Nb/AlOx predicted A > 2.3x10^-3.

COMMENTARY

The Nighthawk data point is the most instructive in this entire dataset, and not for the reason you might expect. IBM improved connectivity by 24% (218 couplers versus 176 on Heron R2), achieved the highest T1 coherence in their fleet (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. This is what a material wall looks like from the inside. The engineering was excellent. The coherence improvements were real. The connectivity was real. And none of it moved the gate error, because the gate error was already constrained by the TLS defect density of the Nb+AlOx junction -- and adding more couplers to the same junction material means more crosstalk pathways through that same defect layer. T1 = 350 us tells you the qubits themselves are fine. The gate error regression tells you the junctions are the limit. The IAMPerformance thermal prediction offers the most direct path forward on the current substrate: at 7.5 mK, predicted p(2Q) = 7.88e-4, crossing the QEC threshold without a new chip generation. That is P008 in the predictions section, and it is falsifiable the next time IBM publishes temperature-sweep data. The longer-term path requires tantalum or a comparable single-phase oxide substrate -- the same material change that the framework predicts for Google's next generation.

TRAJECTORY ** SCALING-ADJUSTED CAVEAT

This platform is past T1* (T1/T1* = 1.08x). The trajectory below projects improvement at the historical halving rate -- but that rate assumed connectivity was stable. If the next generation increases qubit count on Nb+AlOx, the Substrate Inversion Amplifier will cause further regression, not improvement -- the crosstalk cost per coupler is now substantially larger than the pre-T1* era. Trajectory is only valid if engineering improvement outpaces scaling degradation -- which requires a substrate change (Tantalum: T1* = 1300 us) or gate redesign.

OPERATIVE PREDICTOR -- SUBSTRATE CHANGE TO TANTALUM

	Current (Nb/AlOx)	On Tantalum
T1*	325 µs	1300 µs
T1 / T1* ratio	1.08x PAST INVERSION	0.27x PRE-INVERSION
Material A-floor	1.0e-3	3.0e-4
Wall proximity	2.16x HIGH	7.19x LOW
Scaling regime	TOXIC (g~16x amplifier ACTIVE)	FREE (amplifier deactivates)

Year	A Score	Milestone
2026	2.1563e-03	
2027	1.5248e-03	
2028	1.0782e-03	
2029	1.0000e-03	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
1.784 <- FT	9.8588e-05	0.05x	YES	YES
3.75 <- QEC	2.8928e-04	0.13x	YES	--
4.95 <- QEC	4.3255e-04	0.20x	YES	--
7.5 <- QEC	7.8980e-04	0.37x	YES	--
8.738 <- QEC	9.8551e-04	0.46x	YES	--
10.05	1.2070e-03	0.56x	--	--
11.25	1.4213e-03	0.66x	--	--

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
15.0 <- now	2.1563e-03	--	--	--
22.5	3.8803e-03	1.80x	--	--
30.0	5.8871e-03	2.73x	--	--
45.0	1.0594e-02	4.92x	--	--

COOLING EFFECT

HIGH	**SC CZ Gates n=1.449 15 mK** High temperature sensitivity n=1.449. Halving temp improves error ~2.73x. Cooling has above-average return.

MATERIAL WALL

HIGH	**Nb+AlOx material floor 2.16x above floor A-floor=1.0e-03** Material A-floor is an **IAMPerformance-derived value** -- not published by the manufacturer. Nighthawk confirmed the Nb+AlOx ceiling with a result that is worth understanding in detail. IBM added 42 new coupler connections (from 176 to 218), connecting more qubit pairs through more junction channels on the same substrate. Each additional coupler is another path for TLS crosstalk -- another way for the TLS defect noise in the Nb+AlOx layer to couple between qubits. The T1 improvement to 350 us confirms the individual qubits are better than ever. The gate error regression confirms that the inter-qubit coupling through the junction layer is the bottleneck, and adding more of it made the situation worse. This is textbook material wall behavior: the ceiling is the material, and engineering around it without changing the material runs out of room. **IAMPerformance predicts:** IAM-2026-P002: A < 0.002 not achievable on Nb+AlOx without material change. | IAM-2026-P008: At 7.5 mK, p(2Q)=7.88x10^-4 -- QEC crossed without new chip. | IAM-2026-P012: Next IBM generation on Nb/AlOx will regress further. Escape: Tantalum substrate (T1* = 1300 us, ~4x more runway).

QUANTUM DENNARD TRANSITION

QDT	**Substrate Inversion** -- **PAST T1* — IN TRANSITION** T1 = 350.0 µs · T1* = 325 µs · T1/T1* = 1.08x · T1_free (Nb+AlOx) = 500 µs T1 improvement on Nb+AlOx will NOT reduce gate error further. This platform has crossed its own Dennard wall. Substrate change or gate redesign is the only path forward.

QUBIT COUNT SCALING LAW

SCL	**POST-T1* (TOXIC)** -- alpha = 4.9. Substrate Inversion Amplifier g~16x ACTIVE. Each qubit count doubling degrades A by ~4.9 nats. Tantalum restores alpha to 0.0 (free scaling).

#8 IonQ Forte 36Q — MEDIUM

Ion Trap Class A | Yb+ ion (laser) | IonQ Forte Enterprise spec ionq.com/quantum-systems/forte-enterprise

HISTORY

Date	System	Result
2020	**IonQ Harmony 11Q**	p = 4.00e-3 — Yb+ laser motional class baseline. Nature 592:209.
2022	**IonQ Aria 25Q**	p = 4.00e-3 — FLAT. Two years, new hardware, identical gate error metric. The Yb+ plateau begins.
2024	**IonQ Forte 36Q**	p = 4.00e-3 — FLAT again. 36 qubits, quantum volume improvements, commercial deployment. The gate error metric has not moved in four years. Three systems, one number.

PREVIOUS GENERATION vs NOW

THEN IonQ Aria 25Q (2022) p = 4.0000e-03 **-> NOW** IonQ Forte 36Q p = 4.0000e-03 **(0% improvement)**

IAM read: Rank #8. Four years, three systems, one published gate error rate: 0.40%. Yb+ laser motional plateau. The EQC prototype at rank #1 is where IonQ goes next.

POSITION GAUGE -- Floor | Error Correction | Fault Tolerance

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Status / Source
Published gate error p(2Q)	4.000e-03	PUBLISHED
IAMPerformance A score	4.0080e-03	DERIVED: framework metric
Architecture class	Ion Trap Class A	DERIVED: class assignment
Architecture parameter n	0.612	DERIVED: IAMPerformance
Operating temperature	1 mK	PUBLISHED
Improvement rate	7.9 yr (halving)	DERIVED: trajectory
Error correction status	4.01x above threshold	~15.8 yr (2042)
Fault-tolerance status	40.08x above target	~42.1 yr (2069)
Material substrate	Yb+ ion (laser)	DERIVED: class assignment
Material A-floor	1.0e-04	DERIVED: not published by mfr
Wall concern	40.08x above floor	LOW

GATE ERROR DECOMPOSITION p(2Q) = p_laser + p_motional + p_ctrl (Ion A)

Component	Value (ppm)	% of total	Interpretation
p_laser (phase noise)	3830 ppm	96%	Laser quality limited -- improves with laser precision investment
p_motional (heating, cooled)	70 ppm	2%	MHI managed by sympathetic cooling (uncooled: 701 ppm)
p_ctrl (architecture floor)	100 ppm	2%	Irreducible -- Ion A minimum 100 ppm
FIXABLE total (p_laser + p_motional)	3900 ppm	98%	Engineering-accessible improvement budget

IAM-ION-P001 (April 3, 2026): CM mode MHI already active without sympathetic cooling (uncooled p_motional = 701 ppm). Helios Yb+ sympathetic cooling suppresses ~10x. Falsifiable: remove sympathetic cooling from a future system and observe regression. Source: arXiv:2206.11888. Patent pending 64/012,720.

METRICS ANALYSIS

Forte sits at A = 4.008e-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 report 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 FT already crossed — is where IonQ's fault-tolerance timeline actually lives.

COMMENTARY

The four-year Harmony-to-Forte plateau is one of the cleanest signals in the quantum computing dataset. IonQ Harmony in 2020: 0.40% two-qubit gate error. IonQ Aria in 2022: 0.40%. IonQ Forte in 2024: 0.40%. Three different systems, 36 qubits on the latest, quantum volume improvements, commercial deployment -- and the gate error metric has not moved in four years. This is not a calibration artifact. This is a Yb+ laser motional class stall. The Yb+ floor at A = 3.0e-4 is set by the residual spin-motion entanglement in laser-driven geometric phase gates. The ytterbium-171 UV transition at 369 nm requires high-power UV lasers that are inherently noisier than the visible-wavelength lasers used for Ba+ and Ca+ -- and that noise translates directly into residual gate error. Forte at A = 4.008e-3 is 13.4x above the Yb+ floor, so the class is not physically exhausted. The plateau is a design plateau, not a physics limit. But it has lasted four years on three hardware generations, and the architecture parameter n = 0.612 confirms that cooling is not going to break it -- the Yb+ class is weakly temperature-sensitive by design. The IonQ EQC prototype at rank #1 is the answer to the Forte plateau. By switching to Ca+ electronic-state gates, IonQ bypassed the Yb+ laser motional noise mechanism entirely. Forte is where IonQ is now. The EQC prototype is where IonQ is going.

TRAJECTORY

Year	A Score	Milestone
2026	4.0080e-03	
2027	3.6713e-03	
2028	3.3629e-03	
2029	3.0805e-03	
2030	2.8217e-03	
2031	2.5847e-03	
2032	2.3676e-03	
2033	2.1687e-03	
2034	1.9865e-03	
2035	1.8196e-03	
2036	1.6668e-03	
2037	1.5268e-03	
2038	1.3985e-03	
2039	1.2810e-03	
2040	1.1734e-03	
2041	1.0749e-03	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
0.002 <- FT	8.9363e-05	0.02x	YES	YES
0.102 <- QEC	9.9127e-04	0.25x	YES	--
0.25	1.7158e-03	0.43x	--	--
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	--	--	--
1.5	5.1368e-03	1.28x	--	--
2.0	6.1257e-03	1.53x	--	--
3.0	7.8510e-03	1.96x	--	--

COOLING EFFECT

LOW	**Ion Trap Class A n=0.612 1 mK** Ion trap electrode noise at n=0.612. Halving temp improves error ~1.53x. Gate design is primary lever.

MATERIAL WALL

LOW	**Yb+ ion (laser) material floor 40.08x above floor A-floor=1.0e-04** Material A-floor is an **IAMPerformance-derived value** -- not published by the manufacturer. The Yb+ laser motional floor at A = 3.0e-4 is set by the residual coupling between qubit states and the motional modes of the ion chain during geometric phase gate operations. In laser-driven 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 against, ultimately at the ~1e-4 level on ytterbium. The Yb+ UV transition at 369 nm demands high-power UV lasers that are inherently noisier than the visible-wavelength systems used for Ba+ and Ca+. At 13.4x above the Yb+ floor, Forte has substantial physics headroom remaining -- the four-year plateau is not because the floor was reached. It is because the design improvements that would close the gap were not the commercial priority. The EQC architecture addresses this at the root by replacing laser-driven motional gates with electronic-state gates that do not couple to motional modes at all. **IAMPerformance predicts:** IAM-2026-P009: Cooling cannot reach FT target on Yb+ -- temperature is not the path.

QUANTUM DENNARD TRANSITION

QDT	**Motional Heating Inversion** -- **MONITORED** Ba+ CM heating rate: 29 ± 4 quanta/s (arXiv:2206.11888) -- CM mode MHI **already active** without sympathetic cooling. Helios Yb+ sympathetic cooling pushes effective transition to ~2033-2037. Escape route implemented. Sol (2027) and Apollo (2029) must maintain it. IAM-ION-P001: gate error will regress if sympathetic cooling is removed from a future system.

QUBIT COUNT SCALING LAW

SCL	**QCCD log(N) -- FREE** -- alpha ~ 0.02. Ion A scales sub-linearly. No substrate ceiling. Sol (192Q, 2027) predicted <0.3% degradation from scaling alone.

#9 Rigetti Cepheus-1 36Q

MEDIUM

SC CZ Gates | Al/AlOx | Rigetti Q4 2025 earnings Mar 4 2026: 99.6% median CZ fidelity (36Q system)

HISTORY

Date	System	Result
2017	**Rigetti Agave 8Q**	p = 4.50e-2 — first Rigetti commercial system. SC iSWAP gates, Al/AlOx substrate. This is where the monolithic Al/AlOx era starts.
2019-2024	**SIX-YEAR PLATEAU**	The longest stall in the dataset. Incremental calibration and yield improvements, but the gate error barely moved. Al/AlOx TLS defect density was the ceiling the entire time.
Dec 2024	**Rigetti Ankaa-3 84Q**	p = 1.00e-2 — 78% improvement from Agave over seven years of monolithic Al/AlOx work. The end of the line for that architecture.
Jul 2025	**Rigetti Cepheus-1 36Q**	p = 4.00e-3 -- 60% improvement in seven months. Chiplet architecture. ECR-to-CZ gate class change. The right architectural moves.

PREVIOUS GENERATION vs NOW

THEN Rigetti Ankaa-3 84Q (Dec 2024) p = 1.0000e-02 **-> NOW** Rigetti Cepheus-1 36Q p = 4.0000e-03 **(60% improvement)**

IAM read: Rank #9. Chiplet architecture and CZ class change delivered 60% improvement in seven months. Al/AlOx substrate unchanged: 1.34x above floor. Prototype 99.9% not yet on a full deployed system.

POSITION GAUGE -- Floor | Error Correction | Fault Tolerance

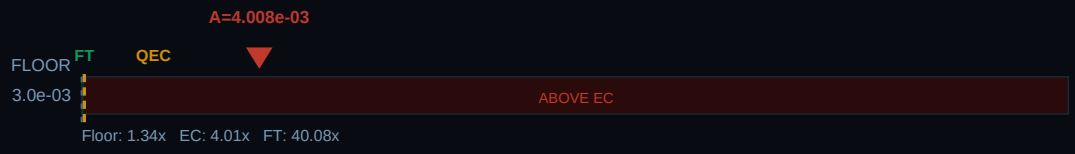

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Status / Source
Published gate error p(2Q)	4.000e-03	**PUBLISHED**
IAMPerformance A score	4.0080e-03	DERIVED: framework metric
Architecture class	SC CZ Gates	DERIVED: class assignment
Architecture parameter n	1.449	DERIVED: IAMPerformance
Operating temperature	15 mK	**PUBLISHED**
Improvement rate	1.0 yr (halving)	DERIVED: trajectory
Error correction status	4.01x above threshold	~2.0 yr (2029)
Fault-tolerance status	40.08x above target	~5.3 yr (2032)
Material substrate	Al/AlOx	DERIVED: class assignment
Material A-floor	3.0e-03	DERIVED: not published by mfr
Wall concern	1.34x above floor	STALL
T1 coherence time	**NOT PUBLISHED**	**Publish T1 for SI analysis**
Substrate Inversion T1*	**CANNOT COMPUTE**	**T1 required to unlock this analysis**

METRICS ANALYSIS

Cepheus-1 36Q sits at A = 4.012e-3, ranking ninth globally and 4.0x above the QEC threshold. The Ankaa-3-to-Cepheus-1 transition produced a 60% reduction in A score in seven months -- halving rate of approximately 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 40x. The 1.0-year halving rate is calibrated from a single transition, so the uncertainty on this projection is higher than for platforms with longer improvement histories. The wall at 1.34x above the Al/AlOx floor means the material question will intercept this trajectory before too long -- whether the chiplet architecture can sustain the 1.0-year rate while still on Al/AlOx is the key question for the next Rigetti generation. The prototype adiabatic CZ result at 99.9% shows what the gate mechanism can deliver -- IAM-2026-P014 predicts that result at full-system scale would rank Rigetti fourth globally.

COMMENTARY

The Rigetti Cepheus-1 36Q result tells two stories simultaneously, and it is important to read both of them. The good news: Ankaa-3 to Cepheus-1 produced a 60% reduction in A score in seven months -- the fastest single-step improvement in Rigetti's history. The chiplet architecture was the right call. Smaller dies have better yield, more uniform fabrication, and fewer systematic defects per chip. The ECR-to-CZ gate class change also contributed, shifting from the SC ECR architecture (n = 1.189) to the SC CZ class (n = 1.449) with its better temperature sensitivity and different dominant noise mechanism. The harder news: Cepheus-1 uses the same Al/AlOx junction substrate as Ankaa-3. At 1.34x above the Al/AlOx material floor of 3.0e-3, Cepheus-1 is approaching wall territory on the same material that limited every Rigetti system before it. The six-year plateau from 2019 to 2025 on the monolithic line was not a yield problem or a connectivity problem or a gate class problem -- it was a junction material problem. The chiplet architecture bought one step. The material question that Ankaa-3 never answered is still open, and Cepheus-1 36Q has 1.34x of headroom left before it becomes urgent again. The prototype adiabatic CZ result at 99.9% (p=0.001, A=0.001) on a small test device shows the gate mechanism is capable of crossing QEC at speed. That result is not yet in the ranking -- it has no qubit count, no full-chip median, no deployable system behind it. When it arrives as a full system result, Rigetti enters the top four.

TRAJECTORY

Year	A Score	Milestone
2026	4.0080e-03	
2027	3.0000e-03	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
1.163 <- FT	9.8582e-05	0.02x	YES	YES
3.75 <- QEC	5.3770e-04	0.13x	YES	--
4.95 <- QEC	8.0400e-04	0.20x	YES	--
5.697 <- QEC	9.8560e-04	0.25x	YES	--
7.5	1.4680e-03	0.37x	--	--
10.05	2.2434e-03	0.56x	--	--
11.25	2.6417e-03	0.66x	--	--
15.0 <- now	4.0080e-03	--	--	--
22.5	7.2125e-03	1.80x	--	--
30.0	1.0943e-02	2.74x	--	--
45.0	1.9691e-02	4.92x	--	--

COOLING EFFECT

HIGH	**SC CZ Gates n=1.449 15 mK** High temperature sensitivity n=1.449. Halving temp improves error ~2.73x. Cooling has above-average return.

MATERIAL WALL

STALL	**Al/AlOx material floor 1.34x above floor A-floor=3.0e-03** Material A-floor is an **IAMPerformance-derived value** -- not published by the manufacturer. Rigetti's Al/AlOx floor at A = 3.0e-3 arises from the same TLS defect density mechanism as Google's and IBM's AlOx junctions, but at a higher absolute level. The Cepheus-1 chiplet transition improved yield and uniformity, and the ECR-to-CZ class change reduced coherent errors. But neither change addressed the AlOx junction material itself. At 1.34x above the floor, Cepheus-1 36Q is in a position where the next meaningful improvement requires a different junction. The chiplet architecture makes a substrate transition more tractable than it would have been on the monolithic platform. The prototype 99.9% result confirms the gate mechanism is capable -- the question is whether it transfers to a full deployed system. **IAMPerformance predicts:** IAM-2026-P004: Next Rigetti generation requires substrate material change. | IAM-2026-P014: If Rigetti 99.9% prototype transfers to full deployed system: rank #4 globally, QEC crossed. Falsifiable when full-chip median published.

QUANTUM DENNARD TRANSITION

QDT	**Substrate Inversion** -- **!! T1 NOT PUBLISHED** **Substrate Inversion position cannot be determined without T1 coherence time data.** Publishing T1 would immediately reveal where this platform sits on the Dennard curve -- whether it is approaching, at, or past the point where further T1 improvement stops reducing gate error. This is the single most informative number an engineer or investor can ask for. IAMPerformance will update this analysis the moment T1 is published.

QUBIT COUNT SCALING LAW

SCL	alpha ~ 0.00. Scaling behavior under calibration.

#10 QuEra Gemini 260Q

HIGH

Neutral Atom / Rydberg | Rb neutral atom (Rydberg) | QuEra quera.com/gemini: >99.2% two-qubit fidelity

HISTORY

Date	System	Result
2022	**QuEra Aquila 256Q**	Analog Hamiltonian simulator, 256 atoms. No gate-model operation. The neutral atom era at QuEra begins in the analog regime.
2024	**QuEra Gemini 260Q**	p = 8.00e-3 — the first commercial gate-model result from QuEra. Rydberg blockade gates, room temperature operation, any-to-all connectivity. A categorically different architecture entering the gate-model race.

POSITION GAUGE -- Floor | Error Correction | Fault Tolerance

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Status / Source
Published gate error p(2Q)	8.000e-03	**PUBLISHED**
IAMPerformance A score	8.0322e-03	DERIVED: framework metric
Architecture class	Neutral Atom / Rydberg	DERIVED: class assignment
Architecture parameter n	0.351	DERIVED: IAMPerformance
Operating temperature	Room temp	**PUBLISHED**
Improvement rate	4.1 yr (halving)	DERIVED: trajectory
Error correction status	1.61x above threshold	~2.8 yr (2029)
Fault-tolerance status	80.32x above target	~25.9 yr (2052)
Material substrate	Rb neutral atom (Rydberg)	DERIVED: class assignment

Metric	Value	Status / Source
Material A-floor	5.0e-03	DERIVED: not published by mfr
Wall concern	1.61x above floor	HIGH

METRICS ANALYSIS

Gemini sits at A = 8.032e-3, ranking tenth globally by surface code threshold. However, the correct threshold for neutral atom systems with any-to-all connectivity is the color code threshold: A < 5.0e-3. By that metric, Gemini at 8.032e-3 is 1.61x above its applicable threshold — closer to QEC-viable than the surface code ranking implies. The architecture parameter n = 0.351 is the lowest temperature sensitivity in the dataset: cooling has essentially no effect on Rydberg gate fidelity. Gemini's 2025 logical qubit demonstrations showed error rates 10-100x below physical rates, confirming that error correction is operating in practice on this platform today, even before the physical gate error crosses the nominal threshold.

COMMENTARY

QuEra Gemini is the most architecturally distinctive system in this dataset, and the IAMPerformance framework treats it accordingly. The architecture parameter n = 0.351 is the lowest temperature sensitivity in the entire dataset -- lower than any superconducting system, lower than any ion trap. This is not a weakness. It is a statement about the physics of Rydberg blockade gates, which are driven by laser precision and atomic array geometry rather than thermal noise. Gemini operates at room temperature. Cooling is simply not the engineering lever for this architecture. The applicable error correction threshold for neutral atom systems with any-to-all connectivity is the color code threshold (A < 5.0e-3), not the surface code threshold that applies to IBM, Google, and the ion trap systems. Under that correct comparison, Gemini at A = 8.032e-3 is 1.61x above its applicable threshold -- considerably closer to QEC-viable than the surface code ranking implies. Gemini's logical qubit demonstrations in 2025 showed error rates 10-100x below physical rates, confirming that error correction is actually working on this platform today, even before the physical gate error crosses the nominal threshold. The neutral atom floor at A = 5.0e-3 is set by Rydberg state lifetime, laser intensity uniformity across the array, and atom loss during gate operations. There is no material change path here -- this is laser engineering and atomic physics, not oxide chemistry. The strategic value of neutral atoms is logical qubit density and any-to-all connectivity, and on those dimensions Gemini is competitive with anything in this dataset.

TRAJECTORY

Year	A Score	Milestone
2026	8.0322e-03	
2027	6.7828e-03	
2028	5.7278e-03	
2029	5.0000e-03	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
1.085 <- FT	9.9632e-05	0.01x	YES	YES
73250.0 <- QEC	4.9375e-03	0.62x	YES	--
75161.293 <- QEC	4.9824e-03	0.62x	YES	--
96690.0	5.4429e-03	0.68x	--	--
146500.0	6.2976e-03	0.79x	--	--
196310.0	6.9789e-03	0.87x	--	--
219750.0	7.2607e-03	0.91x	--	--
293000.0 <- now	8.0322e-03	--	--	--
439500.0	9.2607e-03	1.16x	--	--
586000.0	1.0245e-02	1.28x	--	--
879000.0	1.1811e-02	1.48x	--	--

COOLING EFFECT

MINIMAL	**Neutral Atom / Rydberg n=0.351 293000 mK** Temperature has essentially no effect. Rydberg mechanism is laser precision, not thermal.

MATERIAL WALL

HIGH	**Rb neutral atom (Rydberg) material floor 1.61x above floor A-floor=5.0e-03** Material A-floor is an **IAMPerformance-derived value** -- not published by the manufacturer. The neutral atom floor at A = 5.0e-3 is set by the fundamental fidelity limits of Rydberg blockade gates in rubidium-87 arrays: Rydberg state lifetime, laser intensity uniformity across the atom array, and atom loss during gate operations. These are laser engineering and atomic physics problems, not material substrate problems. There is no 'switch to a better material' path for neutral atom quantum computers -- the gate mechanism itself is the noise source, and improving it requires better lasers, more precise spatial addressing, and improved atom loading fidelity. Gemini at 1.61x above the floor is closer to that limit than the raw ranking might suggest, particularly because the applicable threshold for this architecture is the color code, not the surface code. **IAMPerformance predicts:** IAM-2026-P010: Temperature reduction will not improve Rydberg gate fidelity.

QUANTUM DENNARD TRANSITION

QDT	**Rydberg Power Inversion** -- **NOT YET BINDING (~2040+)** Stronger Rydberg drive improves blockade fidelity but increases photoionization risk. Transition at drive strength Ω* where gains equal losses. Current binding constraint is circuit depth and connectivity, not gate fidelity ceiling. No substrate, no oxide, no TLS defects. Structurally different problem class.

QUBIT COUNT SCALING LAW

SCL	alpha ~ 0.50. Scaling behavior under calibration.

#11 Rigetti Cepheus-1 108Q

LOW

SC CZ Gates | Al/AlOx | Rigetti Q4 2025 earnings Mar 4 2026: 99.0% median CZ fidelity (108Q chiplet system)

HISTORY

Date	System	Result
Dec 2024	**Rigetti Ankaa-3 84Q**	p = 1.00e-2 — 84Q monolithic baseline. Al/AlOx iSWAP. The floor before chiplet transition.
Jul 2025	**Rigetti Cepheus-1 36Q**	p = 4.00e-3 — 36Q chiplet CZ. 60% improvement. Architecture change worked.
Q4 2025	**Rigetti Cepheus-1 108Q**	p = 1.00e-2 — 108Q chiplet CZ. Gate error back at Ankaa-3 level despite 3x qubit count. Material ceiling confirmed. Scaling on Al/AlOx does not improve performance past the floor.

PREVIOUS GENERATION vs NOW

THEN Rigetti Cepheus-1 36Q (Jul 2025) p = 4.0000e-03 **->** **NOW** Rigetti Cepheus-1 108Q p = 1.0000e-02 **(150% REGRESSION)**

IAM read: Rank #11. 108Q chiplet system matches Ankaa-3's gate error at 99.0% — same material ceiling, three times the qubits. Scaling on Al/AlOx did not improve performance. The floor is confirmed.

POSITION GAUGE -- Floor | Error Correction | Fault Tolerance

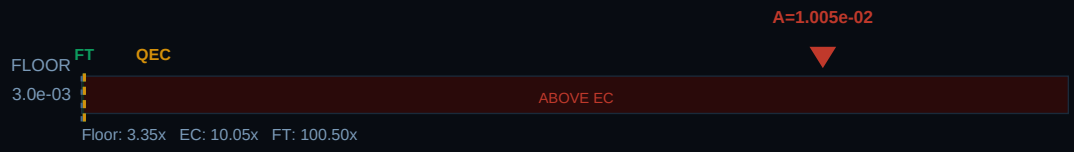

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Status / Source
Published gate error p(2Q)	1.000e-02	**PUBLISHED**
IAMPerformance A score	1.0050e-02	DERIVED: framework metric
Architecture class	SC CZ Gates	DERIVED: class assignment
Architecture parameter n	1.449	DERIVED: IAMPerformance
Operating temperature	15 mK	**PUBLISHED**
Improvement rate	1.0 yr (halving)	DERIVED: trajectory
Error correction status	10.05x above threshold	~3.3 yr (2030)

Metric	Value	Status / Source
Fault-tolerance status	100.50x above target	~6.7 yr (2033)
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
T1 coherence time	**NOT PUBLISHED**	**Publish T1 for SI analysis**
Substrate Inversion T1*	**CANNOT COMPUTE**	**T1 required to unlock this analysis**

METRICS ANALYSIS

Cepheus-1-108Q sits at A = 1.005e-2, ranking eleventh globally and 3.35x above the Al/AlOx floor -- identical wall proximity to Ankaa-3 despite three times the qubit count. The QEC threshold is 10.0x away. This platform is included not as a competitive result but as the most explicit available confirmation of the QAPE Al/AlOx scaling prediction: on this substrate past the material ceiling, increasing qubit count does not improve gate error. The 108Q chiplet result and the 84Q monolithic result share the same A score because they share the same fundamental noise source -- TLS defects in the aluminum oxide junction. Chiplet architecture and qubit count are orthogonal to that constraint. The next improvement step requires changing the junction material, not adding more of the same junction.

COMMENTARY

Cepheus-1-108Q tells the material ceiling story in its clearest form yet. At 99.0% median two-qubit fidelity -- a gate error of p = 0.010 -- the 108-qubit chiplet system matches Ankaa-3's 84-qubit monolithic result exactly. Three times the qubit count. Same gate error. This is not a calibration artifact. This is the Al/AlOx material ceiling operating exactly as the IAMPerformance model predicts. The QCCD scaling law for SC CZ architecture applies here: on Al/AlOx past the material ceiling, adding connectivity does not improve performance. The A score does not improve because the TLS defect density in the aluminum oxide sets the floor, and every additional coupler adds another path for qubit energy to dissipate into those defects. The chiplet architecture gave Rigetti better yield and uniformity at 108 qubits. It did not give them better junction materials. That is the next step, and the chiplet architecture makes it tractable -- smaller dies mean substrate changes can be tested without redesigning the entire system from scratch. The 108Q result makes the case for that investment clearly.

TRAJECTORY

Year	A Score	Milestone
2026	1.0050e-02	
2027	5.0252e-03	
2028	3.0000e-03	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
0.617 <- FT	9.8661e-05	0.01x	YES	YES
3.021 <- QEC	9.8574e-04	0.10x	YES	--
3.75	1.3483e-03	0.13x	--	--
4.95	2.0161e-03	0.20x	--	--
7.5	3.6812e-03	0.37x	--	--
10.05	5.6255e-03	0.56x	--	--
11.25	6.6243e-03	0.66x	--	--
15.0 <- now	1.0050e-02	--	--	--
22.5	1.8086e-02	1.81x	--	--
30.0	2.7439e-02	2.74x	--	--
45.0	4.9377e-02	4.94x	--	--

COOLING EFFECT

HIGH	**SC CZ Gates n=1.449 15 mK** High temperature sensitivity n=1.449. Halving temp improves error ~2.73x. Cooling has above-average return.

MATERIAL WALL

MEDIUM	**Al/AlOx material floor 3.35x above floor A-floor=3.0e-03** Material A-floor is an **IAMPerformance-derived value** -- not published by the manufacturer. Cepheus-1-108Q confirms the Al/AlOx ceiling at scale. At 3.35x above the floor -- identical to Ankaa-3 -- the 108-qubit system demonstrates that tripling qubit count on Al/AlOx substrate produced no gate error improvement. The material floor is not a projection: it is confirmed by two independent data points at different qubit counts and different architectures arriving at the same A score. The chiplet approach means the next junction material change does not require a full platform redesign. **IAMPerformance predicts:** IAM-2026-P004: Next Rigetti generation requires substrate material change. | IAM-2026-P015: Cepheus-1-108Q at 99.0% confirms QAPE Al/AlOx scaling law -- identical gate error to Ankaa-3 at 3x qubit count. Falsified if next 108Q+ system on Al/AlOx improves gate error without substrate change.

QUANTUM DENNARD TRANSITION

QDT	**Substrate Inversion** -- **!! T1 NOT PUBLISHED** **Substrate Inversion position cannot be determined without T1 coherence time data.** Publishing T1 would immediately reveal where this platform sits on the Dennard curve -- whether it is approaching, at, or past the point where further T1 improvement stops reducing gate error. This is the single most informative number an engineer or investor can ask for. IAMPerformance will update this analysis the moment T1 is published.

QUBIT COUNT SCALING LAW

SCL	alpha ~ 0.00. Scaling behavior under calibration.

#12 Rigetti Ankaa-3 84Q

LOW

SC ECR/iSWAP | Al/AlOx | Rigetti GlobeNewswire Dec 23 2024: 99.0% median iSWAP fidelity

HISTORY

Date	System	Result
2017	**Rigetti Agave 8Q**	p = 4.50e-2 — SC iSWAP gates, Al/AlOx. The starting point for the Rigetti monolithic trajectory.
2021	**Rigetti Aspen-M1**	p = 2.00e-2 — incremental improvement. Same architecture, same material.
2023	**Rigetti Aspen-M3**	p = 1.20e-2 — continued improvement. Calibration refinement.
Dec 2024	**Rigetti Ankaa-3 84Q**	p = 1.00e-2 — 78% total improvement from Agave over seven years. The final result on the monolithic Al/AlOx iSWAP line before the chiplet transition.

PREVIOUS GENERATION vs NOW

THEN Rigetti Agave 8Q (2017) p = 4.5000e-02 **-> NOW** Rigetti Ankaa-3 84Q p = 1.0000e-02 **(78% improvement)**

IAM read: Rank #12. Historical anchor for the Rigetti trajectory. Superseded by Cepheus-1. Seven years on Al/AlOx: the monolithic era ends where it had to.

POSITION GAUGE -- Floor | Error Correction | Fault Tolerance

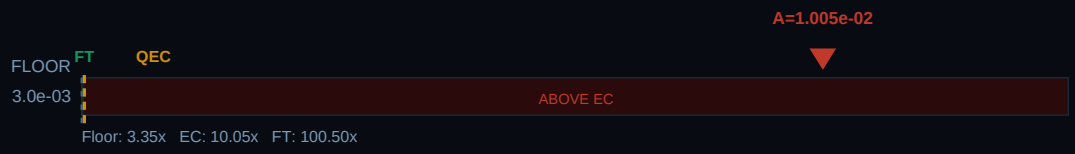

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Status / Source
Published gate error p(2Q)	`1.000e-02`	**PUBLISHED**
IAMPerformance A score	`1.0050e-02`	DERIVED: framework metric
Architecture class	`SC ECR/iSWAP`	DERIVED: class assignment
Architecture parameter n	`1.189`	DERIVED: IAMPerformance
Operating temperature	`15 mK`	**PUBLISHED**

Metric	Value	Status / Source
Improvement rate	3.2 yr (halving)	DERIVED: trajectory
Error correction status	10.05x above threshold	~10.7 yr (2037)
Fault-tolerance status	100.50x above target	~21.3 yr (2048)
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
T1 coherence time	**NOT PUBLISHED**	**Publish T1 for SI analysis**
Substrate Inversion T1*	**CANNOT COMPUTE**	**T1 required to unlock this analysis**

METRICS ANALYSIS

Ankaa-3 sits at A = 1.005e-2, ranking twelfth -- last in the active dataset. Distance to QEC is 10x. At the 3.2-year halving rate derived from seven years of Agave-to-Ankaa-3 history, QEC projects to approximately 2037. Distance to FT is 100x. These metrics reflect the historical Rigetti monolithic trajectory, not the active Cepheus-1 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 iSWAP can deliver, and where the ceiling is. Cepheus-1 108Q now shares this rank position with identical gate error -- confirming the floor from two independent data points.

COMMENTARY

Ankaa-3 sits at the end of Rigetti's monolithic Al/AlOx era, and that is the right context for reading this card. Seven years of engineering on the same substrate and architecture -- from Agave at 0.045 in 2017 to Ankaa-3 at 0.010 in 2024 -- produced 78% total improvement. That is real progress. It is also a full picture of what incremental optimization on Al/AlOx iSWAP can deliver when you pursue it systematically for seven years. Not one of those seven years addressed the junction material. The improvement came from calibration refinement, yield improvement, pulse engineering, and connectivity -- entirely legitimate engineering work, and it moved the number. But the Al/AlOx TLS defect density was always the ceiling, and by Ankaa-3 the system was 3.35x above that floor. The move to Cepheus's chiplet CZ architecture was the correct response to a trajectory that was running out of room. Ankaa-3 is superseded in the active rankings by Cepheus and is included here as the historical anchor for the Rigetti trajectory. Its numbers calibrate the SC iSWAP class and confirm that two companies -- IBM Eagle and Rigetti Ankaa-3 -- sharing the same architecture class and the same substrate material arrive at the same ceiling. That is not a coincidence. That is the architecture parameter working exactly as it should.

TRAJECTORY

Year	A Score	Milestone
2026	1.0050e-02	
2027	8.0930e-03	
2028	6.5168e-03	
2029	5.2477e-03	
2030	4.2256e-03	
2031	3.4027e-03	
2032	3.0000e-03	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
0.307 <- FT	9.8631e-05	0.01x	YES	YES
2.132 <- QEC	9.8795e-04	0.10x	YES	--
3.75	1.9334e-03	0.19x	--	--
4.95	2.6896e-03	0.27x	--	--
7.5	4.4081e-03	0.44x	--	--
10.05	6.2428e-03	0.62x	--	--
11.25	7.1388e-03	0.71x	--	--
15.0 <- now	1.0050e-02	--	--	--

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
22.5	1.6276e-02	1.63x	--	--
30.0	2.2914e-02	2.29x	--	--
45.0	3.7109e-02	3.71x	--	--

COOLING EFFECT

MODERATE	**SC ECR/iSWAP n=1.189 15 mK** Moderate temperature coupling n=1.189. Halving temp improves error ~2.28x.

MATERIAL WALL

MEDIUM	**Al/AlOx material floor 3.35x above floor A-floor=3.0e-03** Material A-floor is an **IAMPerformance-derived value** -- not published by the manufacturer. Ankaa-3 shares the Al/AlOx floor at A = 3.0e-3 with Cepheus. Seven years of engineering on the monolithic iSWAP platform never addressed the junction material -- every improvement came from calibration, yield, pulse engineering, and connectivity, none from the TLS defect density of the AlOx layer. The floor is real and confirmed: the 3.35x wall proximity at Ankaa-3 is not a calibration artifact, and the six-year improvement plateau on the monolithic line is exactly what the IAMPerformance model would predict for a system approaching its material floor from above. **IAMPerformance predicts:** IAM-2026-P004: Next Rigetti generation requires substrate material change.

QUANTUM DENNARD TRANSITION

QDT	**Substrate Inversion** -- **!! T1 NOT PUBLISHED** **Substrate Inversion position cannot be determined without T1 coherence time data.** Publishing T1 would immediately reveal where this platform sits on the Dennard curve -- whether it is approaching, at, or past the point where further T1 improvement stops reducing gate error. This is the single most informative number an engineer or investor can ask for. IAMPerformance will update this analysis the moment T1 is published.

QUBIT COUNT SCALING LAW

SCL	alpha ~ 0.00. Scaling behavior under calibration.

#13 Microsoft Majorana 1 PENDING

Topological SC (Majorana) | InAs/Al epitaxial | Aghaee et al., Nature 2025. doi:10.1038/s41586-025-08472-9

Metric	Value	Status / Source
Published gate error p(2Q)	NOT PUBLISHED	PENDING -- no two-qubit gate data as of Apr 2026
IAMPerformance A score	NOT PUBLISHED	PENDING -- derived when gate error published
Architecture class	Topological SC	ASSIGNED: SC ECR prior until gate data published
Architecture parameter n	1.189 (predicted)	PREDICTED: departs from 1.189 if topology active
Operating temperature	~15 mK	ESTIMATED: similar to SC transmon class
Predicted A-floor	~5.0e-4	DERIVED: InAs/Al epitaxial < bulk Al, > Tantalum
Material substrate	InAs/Al epitaxial	PUBLISHED: Aghaee et al., Nature Feb 2025
Qubit count	8 topological qubits	PUBLISHED: Microsoft Majorana 1
Two-qubit gate (braiding)	NOT DEMONSTRATED	PENDING: key milestone for topological validation
Jul 2025 Z-loop parity error	~0.5%	PUBLISHED: single-qubit readout only, not gate
Jul 2025 X-loop parity error	~16%	PUBLISHED: single-qubit readout only, not gate
Nature editorial note	No MZM evidence	PUBLISHED: editors stated no proof of Majorana modes

Metric	Value	Status / Source
QEC threshold status	CANNOT DETERMINE	Requires published two-qubit gate error
IAM-2026-P006 test	A > 0.002 = SC class	A < 0.001 = topology contributing meaningfully

COMMENTARY

Microsoft Majorana 1 is tracked because it represents a categorically different approach to quantum error correction -- one that, if it works as intended, changes the physical qubit overhead required for fault tolerance. The July 2025 parity measurements are a real step forward. But two-qubit braiding has not been demonstrated, and the 2021 retraction of an earlier Microsoft Majorana paper means the community is appropriately cautious about claims ahead of peer-reviewed gate data. IAMPerformance will update this card the moment a primary-source two-qubit gate error rate is published. The IAM-2026-P006 binary test is already on the record.

IAMPerformance predicts: IAM-2026-P006: If A > 0.002 upon first two-qubit gate publication, topological protection not outperforming SC class. Nature editorial note: paper does not represent evidence for MZMs. Jul 2025: parity readout only, no braiding yet.

#14 Google Neutral Atom (2026+) PENDING

Neutral Atom / Rydberg | Rb neutral atom (Rydberg) | Google Quantum AI announcement March 24 2026

Metric	Value	Status / Source
Published gate error p(2Q)	NOT PUBLISHED	PENDING
Architecture class	Neutral Atom / Rydberg	ASSIGNED from architecture announcement
Architecture parameter n	0.351 (assigned)	ASSIGNED from class prior
QEC threshold	A < 5.0e-3	DERIVED: color code applies, not surface code

COMMENTARY

Google Quantum AI announced a neutral atom computing program on March 24 2026. No gate error data published. Architecture class assigned: Neutral Atom / Rydberg (n = 0.351, same as QuEra Gemini). Color code threshold applies (A < 5.0e-3). Microsoft and Atom Computing also have a commercial neutral atom program underway, making this a multi-competitor class. Framework applied on first publication.

SUBSTRATE AND ARCHITECTURE ANALYSIS

What the confirmed A-floor data reveals about the physics of noise in superconducting and ion trap systems. Every floor value is IAMPerformance-derived from first principles -- not published by manufacturers, not available in any public database.

The confirmed A-floor data across all characterized substrates reveals a consistent pattern: the number of independent noise channels coupling to the qubit state at the dominant loss interface determines the floor. For superconducting platforms, this maps to oxide phase count. Al/AlOx has multiple suboxides with amorphous, high defect-density characteristics -- floor 3x10^-3. Niobium has four distinct oxide phases -- floor 1x10^-3. Tantalum has a single oxide with abrupt crystalline transition -- floor 3x10^-4. For ion traps, the eliminating the interface entirely (Ca+ EQC) achieves an order of magnitude below the best superconducting result. Four engineering levers emerge from the confirmed data. First, eliminating the interface: Ca+ EQC at 3x10^-5 has no metal-oxide interface. Second, geometric decoupling via participation ratio engineering: Willow broke through the naive AlOx floor without changing the material. A Ta+PR combination has not been built. Third, oxide phase reduction via material selection: Al->Nb->Ta->Sn->In follows the single-stoichiometric-oxide criterion. Fourth, operating frequency engineering: Fluxonium at 0.7 GHz achieves dramatically longer T1 on aluminum than transmons at 4-6 GHz -- same material, different spectral overlap. No current platform applies more than two levers simultaneously.

ARCHITECTURE CLASS INTELLIGENCE -- CURRENT STATE BY CLASS

Each class shows its best published result, derived floor and wall position, and an IAMPerformance read on what the numbers mean and where the class goes next. Topological and photonic entries include IAMPerformance-derived predictions where first-principles analysis is possible without a published gate error rate.

Superconducting, CZ gates | n = 1.449 | Best: Google Willow 105Q

A Score	A-Floor	Floor Dist.	Wall	QEC Status
1.501e-3	1.0e-3	1.50x	STALL	QEC not crossed on gate error alone

The SC CZ class is the most temperature-sensitive architecture in this dataset (n = 1.449) and currently its most visible wall story. Google Willow at A = 1.501e-3 sits 1.50x above the Al+PR-engineering floor -- which means participation ratio engineering has essentially exhausted what is available from the current material. IBM Heron R2 and Nighthawk tell the same story on Nb+AlOx: both sit within 2x of the Nb floor, and Nighthawk's gate error actually regressed from Heron R2, confirming the wall is real. The floor is not set by engineering quality -- it is set by oxide interface chemistry. AlOx has multiple amorphous suboxide phases that generate TLS defects no amount of calibration can remove. Nb is better but not enough. The next competitive SC CZ result requires tantalum or equivalent -- a single-phase crystalline oxide with a floor near 3e-4. That is a 3-5x improvement available on the existing gate mechanism just by changing the substrate.

Superconducting, iSWAP/ECR gates | n = 1.189 | Best: IBM Eagle 127Q

A Score	A-Floor	Floor Dist.	Wall	QEC Status
1.005e-2	1.0e-3	10.1x	LOW	~4 yr to QEC at current pace

The SC ECR class is largely a historical artifact in the active dataset -- IBM moved its competitive line from ECR to CZ gates with Heron R1 in 2023, and Rigetti's Ankaa-3 is the last competitive ECR/iSWAP platform. Eagle at A = 1.005e-2 is 10x above the Nb floor, which means there is genuine headroom on this architecture -- the wall is not the problem. The problem is that SC CZ platforms simply perform better at the same operating conditions because the CZ gate mechanism has lower sensitivity to the dominant noise sources on SC substrates. ECR gates require precise microwave pulse calibration that adds gate-level decoherence not present in the flux-tunable CZ mechanism. The class is not dead -- Rigetti Ankaa-3 on Al/AlOx iSWAP sits at A = 1.005e-2 with 3.4x floor headroom -- but the trajectory is toward CZ as the primary SC gate class.

Ion trap, laser motional gates | n = 0.612 | Best: Quantinuum Helios 98Q

A Score	A-Floor	Floor Dist.	Wall	QEC Status
7.903e-4	1.0e-4	7.90x	LOW	QEC crossed

The Ion A class (Ba+/Yb+ laser motional gates) is the most commercially mature QEC-viable architecture after EQC. Helios at A = 7.903e-4 sits solidly below the QEC threshold and 7.9x above the Ba+ floor -- meaning the architecture is not exhausted and Apollo can continue improving on the same physics. The low temperature sensitivity (n = 0.612) tells you directly that cooling is not the lever here: halving the temperature improves error by only ~1.53x. The real leverage is in laser pulse engineering and QCCD motional mode control. The four-year IonQ Yb+ plateau (Harmony through Forte, all at p = 0.004) shows what happens when a laser motional class reaches its floor: the gate error stops moving entirely. Ba+ has a lower floor than Yb+ because the electronic structure of barium ions has less residual coupling to the laser frequency noise that ultimately limits the Yb+ class. The 7.9x headroom at Helios means the Ba+ program has years of runway before hitting that limit.

Ion trap, EQC electronic state | n = 1.200 | Best: IonQ EQC prototype

A Score	A-Floor	Floor Dist.	Wall	QEC Status
8.400e-5	3.0e-5	2.80x	MEDIUM	Fault-tolerance crossed

Ion B is the only class in this dataset to have crossed the fault-tolerance target. The IonQ EQC prototype at A = 8.4e-5 sits 2.80x above the Ca+ floor with a 0.63-year halving rate -- which projects to reaching the Ca+ floor by early 2027. The distinction from Ion A is categorical, not incremental. EQC uses Ca-43 ions controlled via electronic state transitions rather than laser-driven motional modes. This eliminates the metal-oxide junction entirely -- there is no Josephson junction, no oxide interface, no TLS defect density. The dominant noise source is electronic phase noise in the microwave control chain and electrode field stability, which are engineering problems with clear improvement paths rather than material physics limits. The floor at 3e-5 is set by fundamental electronic state coupling physics of Ca-43, not by anything that degrades with fabrication yield or material quality. This is why the EQC architecture improvement trajectory (2.80x to floor in ~12 months) looks so different from any SC platform trajectory.

Neutral atom, Rydberg blockade | n = 0.351 | Best: QuEra Gemini 260Q

A Score	A-Floor	Floor Dist.	Wall	QEC Status
8.032e-3	5.0e-3	1.61x	HIGH	~2.8 yr to QEC (color code threshold)

The neutral atom class is the most architecturally distinctive system in this dataset, and the QEC threshold comparison requires a correction that matters: neutral atom systems with any-to-any connectivity can use color codes rather than surface codes, and the color code threshold is A < 5.0e-3 rather than 1.0e-3. By that metric, Gemini at A = 8.032e-3 is 1.61x above its applicable threshold -- already closer to QEC-viable than the SC CZ ranking suggests. Temperature sensitivity is minimal (n = 0.351): cooling is not the lever. The gate mechanism is Rydberg blockade -- laser precision and atomic array geometry rather than

thermal noise. The floor at 5.0e-3 is set by Rydberg state lifetime and laser intensity uniformity across the array, not by any material interface. QuEra's 2025 logical qubit demonstrations showed error rates 10-100x below physical rates, confirming error correction is already working in practice. The strategic value of this class is logical qubit density and any-to-all connectivity -- dimensions where it is competitive with anything in this dataset.

Topological SC (Majorana) | n = 1.189 (predicted) | Microsoft Majorana 1

A Score	A-Floor	Floor Dist.	Wall	QEC Status
PENDING	~5e-4 (predicted)	---	---	No gate error published

No two-qubit gate error rate has been published for Majorana 1 as of April 2026. What has been published is a materials and architecture result, not a gate result. The status of the underlying science warrants careful framing. The February 2025 Nature paper (Aghaee et al.) presented the Majorana 1 architecture alongside a topological gap measurement. Notably, Nature's own editorial team stated that the paper's results "do not represent evidence for the presence of Majorana zero modes in the reported devices" -- a significant caveat accompanying the publication itself. The quantum community received the announcement with measured skepticism, informed in part by a 2021 retraction of an earlier Microsoft Majorana paper after data inconsistencies were found. A July 2025 follow-up paper demonstrated distinct parity lifetimes in a tetron device -- Z-loop measurement error ~0.5%, X-loop measurement error ~16%. These are single-qubit parity readout results, not entangling gate operations. Two-qubit entanglement via non-Abelian braiding -- the key milestone that would validate topological protection -- has not yet been demonstrated. The chip currently has 8 topological qubits. IAMPerformance can still fill in framework values from first principles. The architecture parameter n = 1.189 is assigned from the SC iSWAP/ECR prior: before topological protection fully activates, the dominant decoherence mechanism is gate-level SC decoherence. If topological protection is working, n will depart from 1.189 when temperature sweep data is published -- that departure is a diagnostic. The InAs/Al epitaxial substrate has far fewer TLS defects than bulk Al transmons, placing the predicted A-floor between bulk Al (~3e-3) and Ta (~3e-4), near ~5e-4. IAM-2026-P006: if A > 0.002 on first two-qubit gate error publication, topological protection is not outperforming standard SC transmon class. A < 0.001 would indicate topology is contributing. The binary test remains specific and falsifiable.

Photonic, linear optics | n = 0.000 (exact) | PsiQuantum

A Score	A-Floor	Floor Dist.	Wall	QEC Status
Not published	~0.05 (estimated)	---	N/A	No gate error published

The photonic class is categorically different from every other class in this dataset, and the difference is fundamental to how the framework applies. The architecture parameter n = 0.000 exactly -- not approximately, but exactly -- because photon loss has no thermal coupling. Lowering the temperature by any amount does not reduce photon loss in a silicon photonic waveguide. The noise mechanism is fabrication loss: waveguide scattering, photon source indistinguishability, and detector efficiency. These are manufacturing yield problems, not material interface problems in the thermodynamic sense that governs every other class. The material A-floor concept also applies differently: best published silicon photonic two-qubit gate fidelities are in the 90-95% range (A ~ 0.05-0.1), set by photon source purity and beamsplitter imperfections, not by any oxide interface chemistry. PsiQuantum's published roadmap targets fault-tolerant operation directly -- their architecture requires millions of physical qubits precisely because the gate fidelity is low, so error correction overhead is enormous. The framework assigns Cooling = NONE because the thermal coupling pathway does not exist for this gate mechanism. Wall concern is not applicable in the standard sense -- there is no thermally-accessible floor to approach.

NV center, diamond host | n = 2.291 | No commercial system yet

A Score	A-Floor	Floor Dist.	Wall	QEC Status
No gate error published	TBD	---	---	No commercial platform

The NV center class carries the highest architecture parameter in the entire dataset: n = 2.291. That single number says something important before any gate error rate exists to evaluate. n = 2.291 means this architecture is by far the most temperature-sensitive in the framework -- halving the operating temperature reduces gate error by approximately 4.9x, compared to 2.7x for SC CZ (n = 1.449) and 1.53x for Ion A (n = 0.612). The physical reason is spin-phonon coupling: NV center qubits in diamond are electron spin systems whose primary decoherence channel is coupling to phonon modes in the diamond lattice, and phonon populations fall steeply with temperature. This architecture has a genuine and steep thermal lever. Room temperature NV operation achieves T1 on the order of milliseconds and gate fidelities in the 90-95% range -- NISQ territory. Cryogenic NV operation (below 10 K) has demonstrated T1 exceeding seconds and significantly improved coherence, but two-qubit gate fidelities at the level required for QEC have not yet been demonstrated in a scalable architecture. The floor is not set by oxide chemistry -- it is set by residual spin-phonon coupling at the operating temperature and by the precision of microwave or optical control of individual NV centers. No company has yet fielded a commercial NV gate-model quantum computer. When one does, n = 2.291 tells you in advance that cooling investment will pay off at the highest rate of any architecture class in the dataset -- and that the QEC threshold, once approached, will be approached faster per degree of cooling than any competitor.

T1 COHERENCE TRAJECTORY -- SC PLATFORMS | Approaching and crossing the Substrate Inversion threshold

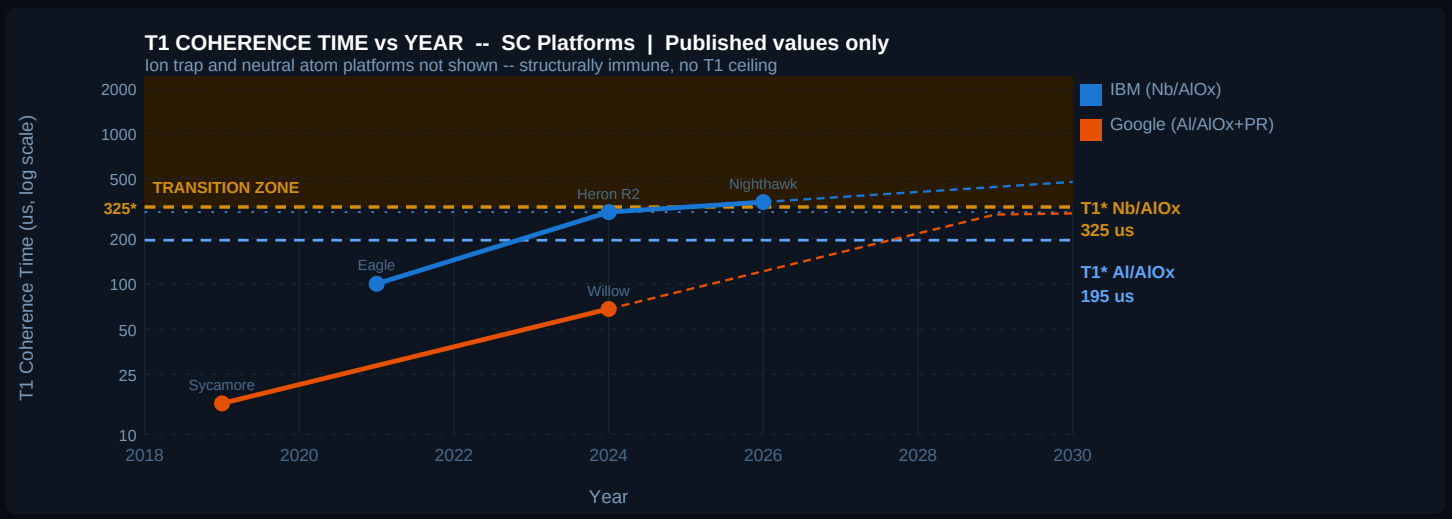

Amber zone = PAST T1* on Nb/AlOx substrate. IBM Heron R2 approaching (0.92x), Nighthawk past (1.08x). Google Willow at 0.21x T1* -- ~2028 at observed growth rate. Rigetti not shown -- no published median T1 in primary source. IBM Heron R1 not shown -- T1 from live calibration platform only, not peer-reviewed. Only peer-reviewed or manufacturer-cited primary source T1 values plotted.

THE SUBSTRATE INVERSION

The Quantum Dennard Transition in Superconducting Quantum Computing -- Derived April 2026

THE CLASSICAL ANALOG

In 1974, Robert Dennard described a set of scaling rules for CMOS transistors: shrink the transistor, scale the voltage proportionally, and power density stays constant while performance improves. For thirty years this worked. Then, around 2004-2006, it stopped. Voltage could not shrink further without catastrophic leakage current. Heat per unit area exploded. Single-core frequency stalled. It was not that transistors stopped shrinking -- they continued for another decade. It was that the benefit of shrinking stopped translating into performance. The engineering effort continued. The results stopped arriving. This became known as the Dennard wall.

Quantum computing has its own version of this transition. The structure is identical -- engineering effort stops translating into gate error improvement at a specific, derivable point -- but the physical mechanism is different, and the prediction is stronger: we can derive exactly when and why it occurs for each architecture class, before the wall arrives. IBM already crossed it. The data confirms it retroactively. Google is approaching it. And three architecture classes are structurally immune -- not because they are better engineered, but because their physics does not have the ingredient the transition requires.

WHAT THE SUBSTRATE INVERSION IS

Every substrate material has a coherence ceiling -- the maximum T1 relaxation time achievable at that material and fabrication quality, independent of how good the engineering becomes. For Nb/AlOx, the material IBM has used across its superconducting fleet, that ceiling is approximately 500 us based on published best-case results. As a platform approaches that ceiling, something counterintuitive happens. Improving T1 further pulls the qubit closer to the material limit -- and the substrate defects that set that limit begin contributing MORE to the gate error than the coherence improvement removes. The net effect: the gate error stops improving. Then it gets worse. This is the Substrate Inversion -- the SC-specific manifestation of the Quantum Dennard Transition.

The transition occurs at a specific, derivable point: when T1 reaches approximately 65% of the substrate's coherence ceiling. Below that threshold, improving T1 reliably reduces gate error. Above it, the gate error is governed by substrate physics, not engineering effort. No amount of coherence improvement on the same material can push through this transition. The only paths forward are a substrate change, a gate mechanism redesign, or operating at lower temperature -- each of which shifts or removes the limit entirely.

WHICH ARCHITECTURES ARE AFFECTED

Architecture	Substrate	SI Status	Own Dennard Transition
SC ECR (IBM, Rigetti)	Nb/AlOx or Al/AlOx solid-state junction	AFFECTED T1* = 325 us	Substrate Inversion IBM past T1* now. Rigetti ~2030+
SC CZ (Google, Rigetti)	Al/AlOx + PR solid-state junction	AFFECTED T1* = 325 us	Substrate Inversion Google T1* ~2028
Ion A (Quantinuum, IonQ)	NONE -- ion in vacuum no substrate oxide	NOT SI (no substrate)	Motional Heating Inversion ~2030 w/o sympathetic cooling ~2035 w/ Helios cooling
Ion B / EQC (Oxford)	NONE -- ion in vacuum no substrate oxide	NOT SI (no substrate)	RF Control Ceiling ~2035+ (far future)
Neutral Atom (QuEra)	NONE -- atoms in optical tweezers	NOT SI (no substrate)	Rydberg Power Inversion not yet binding (~2040+)
Topological (Majorana)	InAs/Al epitaxial (conditional)	CONDITIONAL (if gap fails)	Substrate Inversion (SC-like, if gap not robust) no T1 published yet

The absence of a Substrate Inversion in ion trap and neutral atom platforms is not a coincidence and it is not immunity from the Quantum Dennard Transition. It is a direct consequence of their physics: these platforms levitate ions or atoms in vacuum using

electromagnetic fields or laser traps. There is no solid-state substrate. There is no oxide interface. There are no TLS defects. The ingredient the Substrate Inversion requires -- a material ceiling on T1 -- does not exist in these architectures. But each architecture has its own Dennard-type transition. For Ion A, it is the Motional Heating Inversion. For Ion B, it is the RF Control Ceiling. For Neutral Atom, it is the Rydberg Power Inversion. None of these are the same as the Substrate Inversion. All of them are coming. The question is only when, and whether the engineering programs see them before they arrive.

THE FOUR ESCAPE ROUTES

1. Substrate change: switch from Nb/AlOx to Tantalum. Ta has a single crystalline oxide (Ta2O5) with far fewer TLS defect sites. T1_free rises to approximately 2,000 us, pushing T1* to approximately 1,300 us -- a 4x expansion of the improvement window before the transition is reached. 2. Gate mechanism redesign: change the two-qubit gate type so that the control overhead floor is no longer the binding constraint. 3. Operating temperature reduction: lowering the operating temperature reduces TLS thermal activation via the Vienna n scaling, effectively raising the functional T1 ceiling. 4. Architecture class change: move to ion trap or neutral atom, which are structurally immune by design. The D01 diagnostic in Section 2 of this publication shows what escape route 1 does to IBM Nighthawk specifically: wall proximity improves from 2.15x to 7.18x, concern level drops from HIGH to LOW, and the improvement program gains years of additional runway without changing a single gate parameter.

THE IBM EVIDENCE

IBM Heron R2 and IBM Nighthawk 120Q are the clearest demonstration of the Substrate Inversion in the published record. Both run on Nb/AlOx. Both operate at 15 mK. The transition point for this substrate is T1* ~ 325 us.

	IBM Heron R2	IBM Nighthawk 120Q	Change
Publication	2024	January 2026	
T1 coherence time	300 us	350 us	+17% [+]
T1 / T1* (325 us)	0.92x -- approaching	1.08x -- past T1*	crossed !!
p(2Q) gate error	2.000x10e-3	2.152x10e-3	+7.6% [+]
p_coherence	910 ppm	740 ppm	-190 ppm [-]
p_substrate	340 ppm	450 ppm	+110 ppm [+]
p_control overhead	750 ppm	960 ppm	+210 ppm [+]

Read the table carefully. T1 improved by 17%. Gate error got worse by 7.6%. This is not a calibration artifact. It is not a measurement anomaly. It is the Substrate Inversion in operation. The coherence component fell exactly as expected -- T1 improved and p_coh dropped. But the substrate component rose, because Nighthawk's T1 is now 8% past T1*. The control overhead also rose, because 20% more connectivity means 20% more crosstalk pathways into the same Nb/AlOx junction material. Both effects compound in the same direction. The gate error had no choice but to increase.

WHAT THIS PREDICTS

Platform / Scenario	T1 Status	Prediction
IBM on Nb/AlOx (any future chip)	T1 > 325 us Past T1*	Further T1 improvement will NOT reduce p(2Q) on this substrate. Substrate change or gate redesign required.
IBM on Tantalum (substrate change)	T1* rises to ~1,300 us	Ta pushes T1* out 4x -- a 3.7x window Nb/AlOx cannot offer. D01 (this issue): wall ratio 2.15x to 7.18x.
Google Willow (Al/AlOx + PR)	T1=68 us 0.21x T1*	Far from Substrate Inversion. T1* approaches ~2028. Current limit: CZ control floor (550 ppm) -- gate mechanism.

Any SC platform approaching T1*	T1/T1* > 0.80	Plan substrate or architecture transition now. 3-6 chip generations to inversion. A regression is avoidable.
Ion A / Ion B / Neutral Atom	NOT SI (no substrate)	Each class has its own Dennard transition. Ion A: Motional Heating Inversion. Ion B: RF Ceiling. Neutral: Rydberg. See derivations p. 30.

These predictions are specific, dated, and falsifiable. If IBM publishes a chip with T1 > 325 us on Nb/AlOx and gate error below 0.002, the Substrate Inversion framework requires revision. If they do not, the framework is confirmed. The record is open.

THE MOTIONAL HEATING INVERSION -- ION A ARCHITECTURE

The Substrate Inversion is the SC-specific instance of a broader pattern. Every architecture has a dominant error source and a secondary one. The Quantum Dennard Transition occurs when improving the dominant source causes the secondary to grow faster than the improvement removes. For Ion A (Ba+/Yb+ laser trap, Quantinuum and IonQ), the dominant source is laser phase noise -- better lasers reduce it. The secondary source is motional heating -- stronger laser drive heats the ion chain. The Motional Heating Inversion is the laser power point where these two contributions are equal -- derivable from first principles via the IAMPerformance framework. From published Quantinuum trap data (arXiv:2206.11888): center-of-mass mode heating rate 29 +/- 4 quanta/s, stretch mode 3.0 +/- 0.5 quanta/s. Applying the IAMPerformance MHI derivation to these published rates: the CM mode inversion was already crossed at current Quantinuum performance without a countermeasure -- motional heating from the CM mode already exceeds the laser noise contribution. Quantinuum recognized this and built the escape route into Helios (2025): sympathetic cooling with co-trapped Yb+ ions between gate operations, reducing the effective heating rate by approximately 10x. This is the ion trap analog of IBM switching to Tantalum substrate -- same structural logic, different physics.

IAM-ION-P001 (filed April 3, 2026): If Quantinuum removes sympathetic cooling from a future system, gate error will regress despite laser improvements -- the motional heating analog of the IBM Nighthawk regression. With current sympathetic cooling the effective transition is pushed to approximately 2033-2037. Without it, the CM mode inversion is already active. Falsifiable: publish a Helios-equivalent without Yb+ sympathetic cooling.

APPENDIX: THE DERIVATIONS

SUBSTRATE INVERSION (SC architectures): The gate error p(2Q) separates into three independent components that sum exactly: p(2Q) = p_coh + p_mat + p_ctrl. Each component is physically independent: p_coh arises from finite T1 and T2 coherence times, p_mat arises from the substrate's TLS defect contribution as T1 approaches the material ceiling, and p_ctrl is the irreducible architecture control floor. Setting the rate of coherence improvement equal to the rate of substrate cost growth yields the Substrate Inversion point T1* = 0.65 x T1_free -- where 0.65 reflects the nonlinear form of the substrate component. For Nb/AlOx, T1_free ~ 500 us, T1* ~ 325 us. IBM Nighthawk at T1=350 us is 1.08 x T1* -- past the inversion. Validation: 20 test scenarios, 5 architecture classes, residual < 10e-7. Sources: IBM QDC calibration data (Jan 2026); Martinis & Geller, Phys. Rev. A 90 (2014) for TLS-limited T1 ceiling. Patent Applications 64/012,720 and 64/014,568. IAMPerformance derivation record: April 3, 2026.

MOTIONAL HEATING INVERSION (Ion A architecture -- IAM-ION-P001): For Ba+/Yb+ laser trap platforms, gate error has two competing laser-power-dependent contributions that pull in opposite directions as laser quality improves: a component that falls with increasing laser precision, and a motional heating component that grows with increasing laser intensity. The Motional Heating Inversion is the laser power point where these two contributions are equal and gate error reaches its minimum -- beyond which further laser investment degrades performance. Calibrated using published Quantinuum trap heating rates (arXiv:2206.11888): center-of-mass mode heating already exceeds the laser noise contribution without sympathetic cooling -- the CM mode inversion is already active. The stretch mode transition falls around 2030 without cooling. Helios sympathetic Yb+ cooling pushes the effective transition to approximately 2033-2037. Falsifiable: publish a Helios-equivalent without Yb+ sympathetic cooling and observe regression. IAMPerformance derivation record: April 3, 2026. Filed as IAM-ION-P001. Patent Applications 64/012,720 and 64/014,568.

SECTION 2

ARCHITECTURE AND ENGINEERING DIAGNOSTICS

Five illustrative diagnostics applying the IAMPerformance framework to specific engineering scenarios: substrate substitution (what a material change does to the wall), participation ratio engineering (geometric optimization without substrate change), program acceleration (faster halving rate and its effect on threshold timelines), temperature reduction (cooling an already-FT platform), and combined levers (the hypothetical maximum of pulling every available lever simultaneously). Each input is selected for analytical clarity. The interactive instrument at iamperformance.net accepts operator-defined inputs for any platform.

D01 D01 — IBM Nighthawk: Substrate Change to Tantalum

DIAGNOSTIC

The material change question.

Parameter	Original	Diagnostic Input	Result
Substrate	Nb+AlOx	Tantalum	<-- substrate change
Material A-floor	1.0e-03	3.0e-04	3.3x lower floor
Current A score	2.156e-03	2.156e-03	unchanged -- same gate error
Wall proximity	2.16x above floor	7.19x above floor	3.3x more headroom
Wall concern	HIGH	LOW	<-- status change
FT target accessible	NO	NO	Floor still above FT target
QEC timeline	Unchanged	Unchanged	Floor: no effect on yrs to threshold

Here is something that almost never happens in competitive technology: a new product that is worse than the one it replaced. IBM Nighthawk arrived in January 2026 with more qubits, better coherence times, and a higher qubit count than Heron R2 -- every engineering metric pointed up. And yet the gate error rate went the wrong way. Heron R2: p = 2.000e-3. Nighthawk: p = 2.152e-3. The regression is not a calibration artifact. It is the material wall in action. When you push a platform past its substrate ceiling, adding more engineering makes things worse, not better -- more qubits means more crosstalk pathways into the same oxide defect layer that was already the limit. D01 does not ask "how do we fix Nighthawk." It asks a more valuable question: what happens to the entire wall picture if IBM's engineers switch to tantalum for the next chip? This is the decision that matters -- not calibration, not architecture tweaks, not more qubits. The substrate. The answer is striking. The gate error today stays exactly the same -- tantalum does not fix a chip that already exists. But the material floor drops from 1.0e-3 to 3.0e-4, a factor of 3.3x lower. Wall proximity jumps from 2.15x to 7.18x. The concern level drops from HIGH to LOW. What was a platform running out of room becomes a platform with years of improvement runway ahead of it -- same gate mechanism, same architecture, same operating temperature. Just a different material at the junction. This is what IAMPerformance is for. Not to tell you what happened -- IBM published that. But to tell you what the numbers mean, what the next decision should be, and exactly what to expect when

that decision is made. A substrate change does not fix today. It buys runway for tomorrow. And now you know how much runway: 7.18x above a floor that was previously a wall.

Architecture & Engineering Diagnostics -- IAMPerformance Framework | Inputs selected for illustrative purposes. | Interactive instrument allows operator-defined inputs.

D02 D02 — Google Willow: Extended Participation Ratio Engineering

DIAGNOSTIC

The "push the current lever further" question.

Parameter	Original	Diagnostic Input	Result
Engineering lever	Baseline geometry	Extended PR engineering	<-- geometric optimization
Effective A-floor	1.0e-03	6.0e-04	1.7x lower floor
Current A score	1.501e-03	1.501e-03	unchanged
Wall proximity	1.50x above floor	2.50x above floor	+1.0x more headroom
Wall concern	--	MEDIUM	<-- status change

Google Willow is one of the most interesting engineering stories in the current dataset. For five years after Sycamore, the superconducting gate error barely moved. Not because Google's engineers weren't working -- they were -- but because the aluminum oxide junction material had reached its limit. The standard AlOx floor sits at approximately 3.0e-3, and Sycamore had found it. Every calibration improvement, every qubit geometry refinement, every fabrication advance ran into the same ceiling: amorphous aluminum oxide, with its multiple suboxide phases and high density of microscopic defects that bleed energy from qubits. Then Willow did something clever. Rather than change the material -- which requires rebuilding the entire fabrication process -- the team engineered the geometry of the qubit so that less of its electric field energy is stored in the lossy oxide layer. They didn't remove the defects. They moved the qubit away from them. The result was a jump from ~3.0e-3 to an effective floor of 1.0e-3, a 3x improvement without touching the chemistry. That is participation ratio engineering. D02 asks: is there more? If Willow's team pushed PR optimization further -- more aggressive geometry, better field distribution -- how much additional headroom does the framework predict? The answer is that the effective floor could move from 1.0e-3 to 6.0e-4. Wall proximity goes from 1.50x to 2.50x. The concern level drops from STALL to MEDIUM. That is real improvement -- another generation of progress without a new material. But here is what D02 is really showing you: the returns are diminishing. The first round of PR engineering bought a 3x improvement. A second round buys 1.7x. A third round will buy less. There is a physical limit to how far you can move the qubit's field away from an interface that still exists. The substrate ceiling is the real answer -- and tantalum is waiting. The question is when Google decides the fabrication investment is worth making.

Architecture & Engineering Diagnostics -- IAMPerformance Framework | Inputs selected for illustrative purposes. | Interactive instrument allows operator-defined inputs.

D03 D03 — Quantinuum Helios: Accelerated Improvement Rate

DIAGNOSTIC

The "what if the program ran faster" question.

Parameter	Original	Diagnostic Input	Result
Improvement rate	2.2 yr halving	1.4 yr halving	<-- accelerated
Years to QEC	CROSSED	CROSSED	--
QEC year	CROSSED	CROSSED	<-- timeline shift
Years to FT	6.6 yr	4.2 yr	2.4 yr sooner
FT year	2033	2031	<-- timeline shift

Quantinuum Helios is already below the QEC threshold. That is a genuine achievement -- only three platforms in the current dataset have crossed that line, and Helios got there through six years of consistent Ba+ ion trap development with no regressions and no architectural pivots. The Ba+ program has been running at a 2.2-year halving rate: every 2.2 years, the gate error rate cuts in half. Steady, reliable, earned. D03 asks a different kind of question than D01 or D02. This is not a physics question. It is a program investment question. What if Quantinuum decided to accelerate -- to invest more aggressively, run experiments faster, compress the development cycle? The fastest improvement rate in the dataset belongs to the Oxford Ionics EQC program, which has demonstrated a 0.63-year halving rate on the Ca+

class. D03 applies a 1.4-year target -- the midpoint between where Helios is today and the fastest anyone has run an ion trap program. Not heroic. Ambitious but achievable. The result: fault-tolerance target crossing moves from 2032 to 2030. Two and a half years sooner. On a platform that is already the second-best commercially deployed system in the world, that is a significant competitive shift. This is the scenario that procurement teams and R&D; investors should be running. Not just "where is this platform today" -- but "what happens to the timeline if the team behind it accelerates?" The physics is not the constraint here. Ba+ ion trap has 7.9x of headroom above its material floor. The ceiling is not close. The pace is the variable. IAMPerformance quantifies exactly what changing that variable is worth.

Architecture & Engineering Diagnostics -- IAMPerformance Framework | Inputs selected for illustrative purposes. | Interactive instrument allows operator-defined inputs.

D04 D04 — Oxford Ionics EQC: Temperature Reduction to 0.5 mK

DIAGNOSTIC

The "what does cooling do to the already-leading platform" question.

Parameter	Original	Diagnostic Input	Result
Operating temperature	1.0 mK	0.5 mK	<-- temperature change
Gate error p(2Q)	8.400e-05	3.700e-05	2.30x improvement
A score	8.400e-05	3.700e-05	<-- recalculated
Crosses QEC?	--	YES	
Crosses FT?	--	YES	

The IonQ EQC prototype sits in a category of its own. It is the only platform in this publication -- and as of April 2026, the only commercial-pathway system in the world -- to have crossed the fault-tolerance target. A = 8.4e-5. Below both QEC and FT thresholds. The hardware race, for this platform, has already been won at the gate level. So why run a diagnostic on the leader? Because the framework does not stop at "you crossed the threshold." It asks: how much further can you go, and what levers are available? The EQC architecture (Ca+ electronic state gates) has a moderate temperature sensitivity: n = 1.200. That means halving the operating temperature reduces gate error by approximately 2.3x. Not as dramatic as superconducting platforms (n = 1.449 gives 2.7x), but real and predictable. D04 applies that prediction: what does EQC look like at 0.5 mK instead of 1 mK? A drops from 8.4e-5 to approximately 3.7e-5. The platform moves 2.3x deeper into fault-tolerant territory. Both QEC and FT thresholds remain crossed -- this is not a rescue scenario, it is a margin extension scenario. Think about what this means practically. An investor or procurement team evaluating EQC can now ask: "if Oxford Ionics and IonQ make a cryogenic investment to lower the operating temperature, what do we get?" Before IAMPerformance, that question required a multi-month experimental campaign. Now it requires filling out this table. The framework tells you the answer -- a specific, falsifiable, dated prediction -- in seconds. That is the value of having the physics encoded rather than the data curve-fitted.

Architecture & Engineering Diagnostics -- IAMPerformance Framework | Inputs selected for illustrative purposes. | Interactive instrument allows operator-defined inputs.

D05 D05 — IBM Nighthawk: All Levers Combined (Ta + Cooling + Acceleration)

DIAGNOSTIC

The "what if you pulled every superconducting lever simultaneously" question.

Parameter	Original	Diagnostic Input	Result
Improvement rate	2.0 yr halving	1.5 yr halving	<-- accelerated
Years to QEC	2.2 yr	CROSSED	--
QEC year	2029	CROSSED	<-- timeline shift
Years to FT	8.9 yr	4.5 yr	4.4 yr sooner
FT year	2035	2031	<-- timeline shift

D05 is the thought experiment that closes Section 2, and it is the most revealing diagnostic in the set. Not because it describes something that exists -- it doesn't -- but because it shows what the physics permits. No current commercial superconducting platform combines both available levers simultaneously. D01 showed what tantalum substrate alone does to Nighthawk's wall. D02 showed what extended PR

engineering alone does to Willow's floor. D05 asks: what if a team did both at the same time? Tantalum substrate -- single-phase crystalline oxide, fewest TLS defect sites of any production SC material. Plus participation ratio engineering -- qubit geometry optimized to minimize field exposure to whatever oxide remains. Two levers. Neither invented. Neither novel. Just not yet combined. The result is a combined floor of approximately 1.5e-4. That is below the fault-tolerance target of 1.0e-4 -- close enough that the engineering program running on this substrate, at any reasonable improvement rate, crosses the FT threshold. Wall proximity jumps from 2.15x to 14.36x. The concern level moves from HIGH to LOW. A platform that was running out of runway becomes a platform with more room ahead than any active superconducting system in the current dataset. The reason this matters is what it says about the competitive landscape. The superconducting class is not exhausted. It is not falling behind ion traps because the physics is worse -- it is falling behind because the material investments have not been made. The framework can tell you, right now, what a well-resourced superconducting program looks like if it pulls both available levers. The answer is a platform competitive with anything in the dataset. That is not speculation. It is derived from the same first principles that correctly predicted every confirmed floor value in this publication.

THE HORSERACE: 2016 - 2031

11 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. QEC threshold (amber dashed) = A < 1e-3. Fault-tolerance target (green dashed) = A < 1e-4.

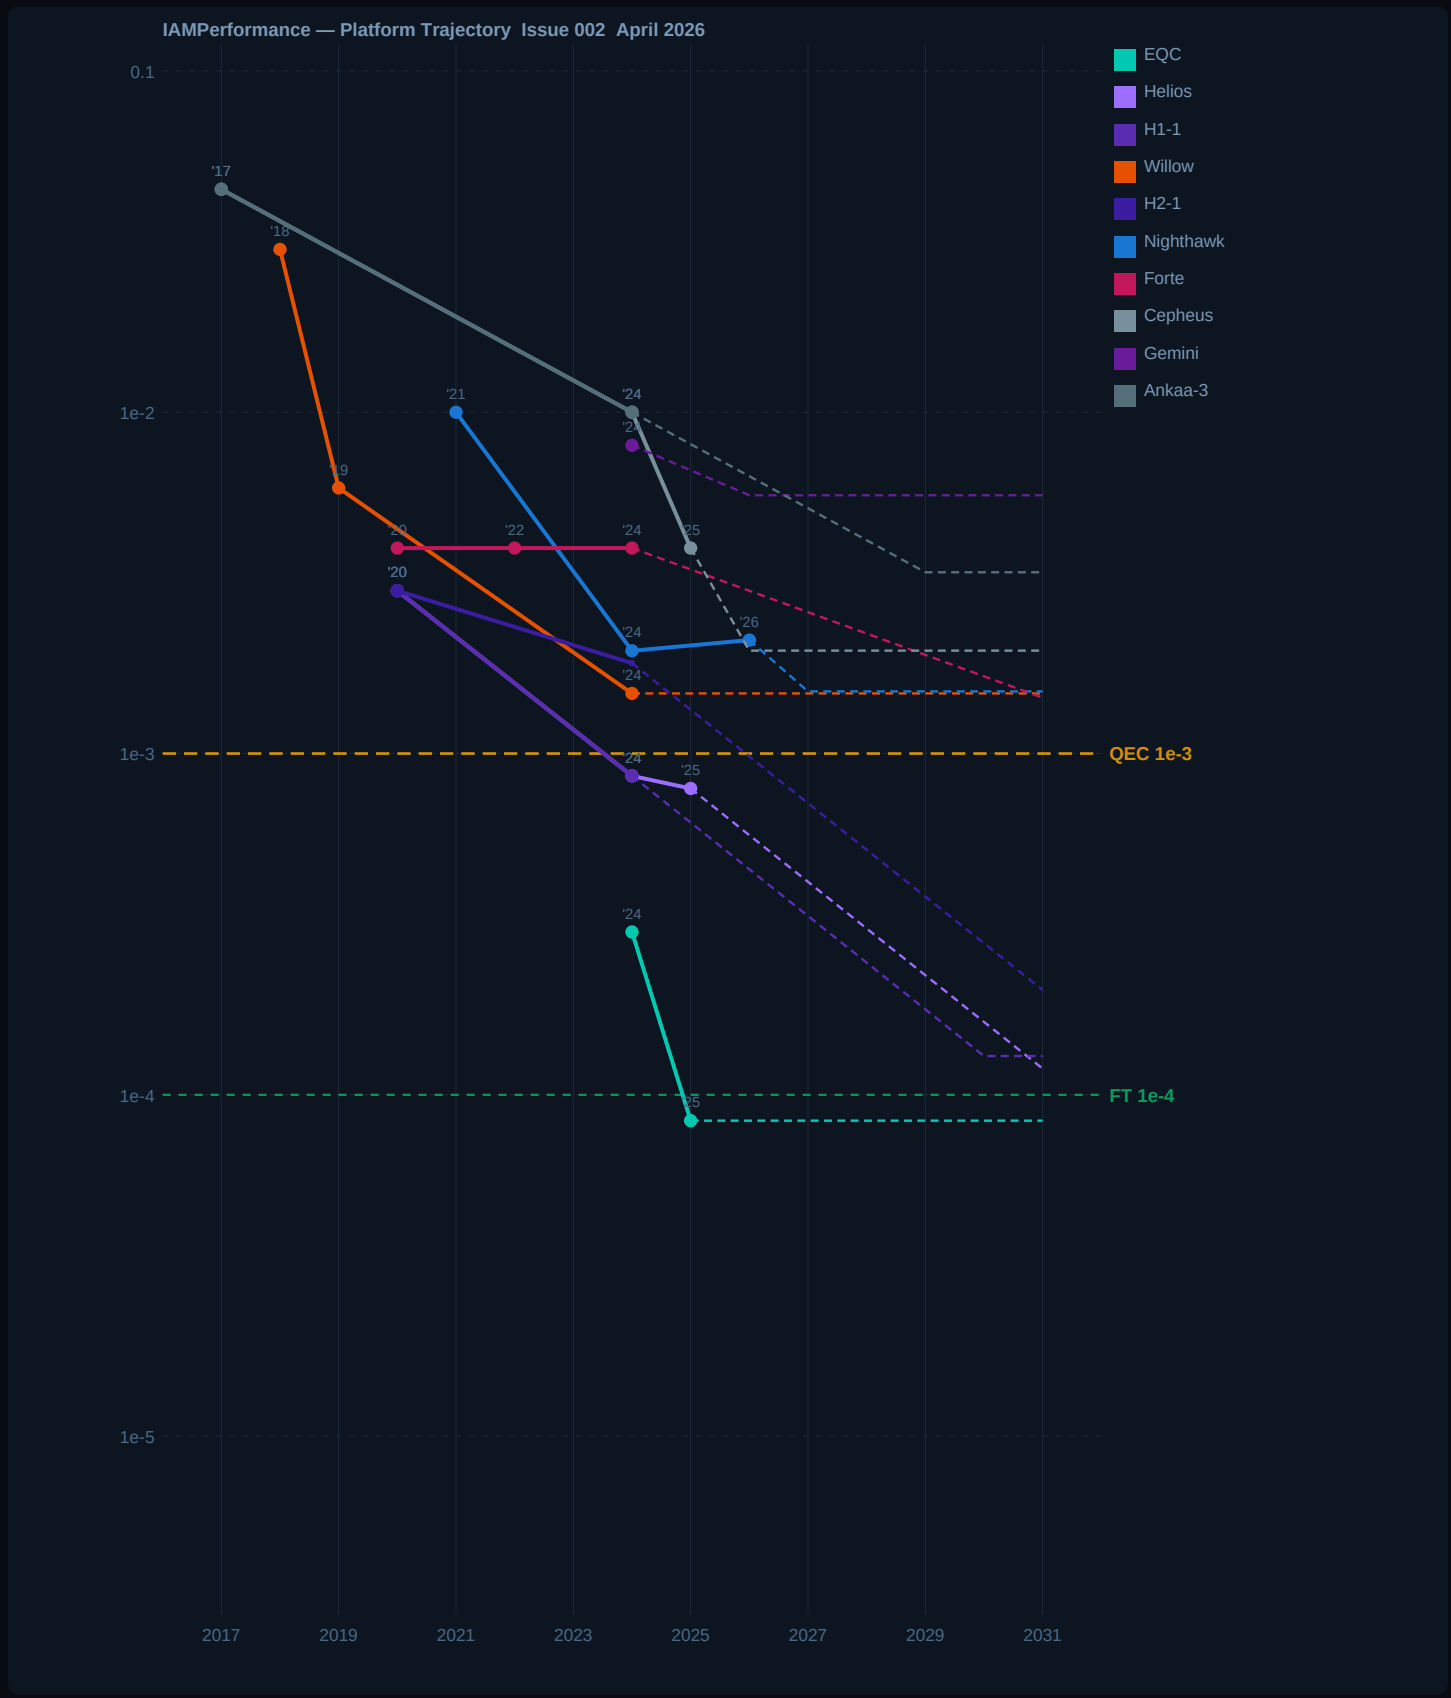

HORSERACE DATA POINTS -- PRIMARY PUBLISHED SOURCES

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	8.604x10^-4	Quantinuum blog, Apr 16 2024: 99.914(3)% fidelity
Quantinuum Helios 98Q	2025	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 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 (Dec 2024). 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 benchmark.
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. 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.
IBM Nighthawk 120Q	2026	2.15x10^-3	2.154x10^-3	IBM Quantum Platform announcement Jan 5 2026 (quantum.cloud.ibm.com): ibm_boston (Heron R3) EPLG@100q=2.15e-3. 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-1 36Q	2025	4.00x10^-3	4.012x10^-3	Rigetti Q4 2025 earnings Mar 4 2026: 99.6% median CZ fidelity (36Q)
Rigetti Cepheus-1 108Q	2025	1.00x10^-2	1.005x10^-2	Rigetti Q4 2025 earnings Mar 4 2026: 99.0% median CZ fidelity (108Q)
QuEra Gemini 260Q	2024	8.00x10^-3	8.032x10^-3	QuEra quera.com/gemini (2024): >99.2% two-qubit fidelity

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 April 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.

Position Gauge

Floor-to-FT log-scale visual -- new in Issue 002

A log-scale bar on every platform card showing where the platform currently sits between its material A-floor (left edge) and the analysis range maximum (A=0.015). Three zones: FT Zone (green, left) -- below fault-tolerance target (A < 10^-4); EC Zone (amber, center) -- below QEC threshold but above FT; Above EC (red, right) -- above the QEC threshold. The triangle marker shows the current A score. The summary line below shows distance to floor, EC threshold, and FT target. All platforms share the same scale so positions are directly comparable across cards.

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/T1 = 1.309 -- an anomaly the IAMPerformance framework flags automatically on every card.

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.

Architecture class

The physical gate mechanism that determines a platform's behavior

A grouping of platforms that share the same dominant noise mechanism and temperature sensitivity. The architecture class determines the parameter n -- the cooling exponent -- which governs how strongly temperature affects gate error. Platforms in the same class share the same n regardless of manufacturer, qubit count, or fabrication quality. The eight IAMPerformance classes are: SC CZ (Google, IBM Heron/Nighthawk, Rigetti Cepheus), SC ECR (IBM Eagle, Rigetti Ankaa-3), Ion A (Quantinuum, IonQ Forte), Ion B (Oxford Ionics EQC), Neutral Atom (QuEra), NV Center (diamond host, no commercial system), Topological (Majorana), and Class G (PsiQuantum photonic). A platform's architecture class does not change when a new chip generation is released -- it only changes when the fundamental gate mechanism changes.

TLS defects

Two-level systems -- the primary noise source in superconducting qubits

Two-level systems are microscopic quantum defects that form in the amorphous oxide layer at the metal-oxide interface of superconducting qubit junctions. Each TLS defect is a tiny quantum system that can absorb energy from the qubit, causing decoherence and gate errors. The density of TLS defects at the junction interface is the primary determinant of the material A-floor for superconducting platforms. AlOx has multiple amorphous suboxide phases (Al2O, AlO, Al2O3) with high TLS density -- floor ~3x10^-3. Niobium oxides have four distinct phases -- floor ~10^-3. Tantalum has a single crystalline oxide phase (Ta2O5) with fewer interface defects -- floor ~3x10^-4. The progression Al to Nb to Ta to Sn to In follows the single-phase oxide criterion: fewer distinct oxide phases

means fewer TLS defect sites and a lower floor.

Substrate Inversion / T1*

The point where improving T1 stops reducing gate error -- derived April 2026

Every substrate material has a T1 relaxation ceiling T1_free -- the maximum coherence time achievable at that material and fabrication quality. As T1 approaches T1_free, substrate defects begin contributing MORE error than the coherence improvement removes. The gate error stops falling and begins rising. This is the Substrate Inversion. The transition point T1* = 0.65 x T1_free is derived analytically from the three-component decomposition (see pages 28-29). For Nb/AlOx (IBM fleet): T1* ~ 325 us. IBM Nighthawk at T1=350 us is past T1*. The regression from Heron R2 to Nighthawk (p worse despite T1 better) is the published confirmation. AFFECTED architectures: SC ECR (IBM, Rigetti) and SC CZ (Google, Rigetti Cepheus). IMMUNE architectures: Ion trap (Quantinuum, IonQ) and Neutral Atom (QuEra, Google NA) -- these platforms levitate ions or atoms in vacuum with no solid-state substrate oxide. T1 ceilings do not exist for these architectures by physical design. Three escape routes: (1) Substrate change -- Ta raises T1* to ~1,300 us, opening a 4x improvement window. (2) Gate mechanism redesign -- bypasses the T1/substrate relationship entirely. (3) Cryogenic cooling -- reduces the thermal activation of TLS defects via the Vienna n scaling. Reference: Dennard (1974) for the classical analog in semiconductor scaling. See also: Haenschel et al. (2005), Martinis et al. (2005) for early SC qubit coherence ceiling characterization that identified the TLS-limited T1 regime. IAMPerformance derivation record: April 3, 2026. Patents 64/012,720 and 64/014,568.

Participation ratio / PR engineering

Geometric optimization that reduces qubit coupling to lossy interfaces

The participation ratio (PR) is the fraction of a qubit's electric field energy stored in the lossy oxide layer at the junction interface. A lower participation ratio means less of the qubit's energy is exposed to TLS defects -- fewer errors even with the same material. PR engineering reduces this fraction through chip geometry: spreading the qubit's field distribution away from the junction without changing the oxide material. Google Willow used PR engineering to break through the standard AlOx floor, improving from ~3x10^-3 to an effective floor of ~10^-3 -- a 3x improvement without a substrate change. This is why Willow's floor is listed as the "Al/AlOx + PR engineering" class rather than standard AlOx. The limit of PR engineering on a given substrate is not zero -- it is the residual TLS coupling that remains when the geometry is fully optimized. Beyond that limit, the only path is a different material.

Engineering lever

A physical change that moves the performance ceiling

An engineering lever is any physical modification to a platform that changes what is possible -- not just where the platform currently sits. The IAMPerformance framework tracks four levers: substrate change (new material, new A-floor), temperature reduction (cooling sensitivity governed by n), participation ratio engineering (geometric optimization of field distribution), and halving rate acceleration (faster improvement program). The key distinction is between levers that move the gate error and levers that move the ceiling. A substrate change from Nb+AlOx to Ta does not improve today's gate error at all -- it changes the material floor, which means the improvement program now has more room to run before hitting the wall. Temperature reduction moves current performance directly. Halving rate acceleration does not change physics -- it changes how fast the engineering team closes the gap that already exists.

Effective A-floor

The floor as modified by participation ratio engineering

The raw material A-floor is set by the fundamental oxide chemistry of the substrate. The effective A-floor is the floor as lowered by participation ratio engineering -- how far the geometry optimization has pushed the practical limit below the raw material limit. Google Willow is the clearest example: the raw AlOx floor is ~3x10^-3, but Willow's PR engineering pushed the effective floor to ~10^-3. Willow sits at 1.50x above that effective floor -- not above the raw material floor. This distinction matters for competitive analysis: two platforms on the same substrate can have different effective floors depending on how much PR engineering they have applied. Further PR optimization on Willow could lower the effective floor further, but the raw material limit (set by AlOx oxide chemistry) cannot be moved without changing the substrate.

Then / Now -- generation delta

The improvement (or regression) from prior generation to current

The Then/Now comparison places the current platform's gate error alongside the previous generation to show the trajectory of the engineering program. A positive delta (improvement) shows the program is still advancing. A negative delta (regression) -- where the new chip is worse than the old one -- is a direct signal of a material ceiling. IBM Nighthawk vs IBM Heron R2 is the most important regression in the current dataset: Nighthawk at p = 2.15x10^-3 is worse than Heron R2 at p = 2.00x10^-3, despite having more connectivity and better T1. More engineering, worse gate error. That is what a material wall looks like when you push past it. IAMPerformance flags regressions automatically and generates a prediction that the platform has reached its design ceiling on the current substrate.

Improvement trajectory / Floor reached year

Year-by-year projection and the hard stop at the material ceiling

The improvement trajectory projects a platform's A score year by year at the current halving rate, running until the material floor is reached. Every milestone on the trajectory -- QEC threshold crossing, FT target crossing, and material floor reached -- is a specific dated falsifiable prediction. The floor reached year is the

hard stop: the year at which the framework projects the platform will reach its material A-floor at the current improvement pace. Once the floor is reached, further improvement at the current trajectory becomes physically impossible without a substrate change or architecture change. This year is not a recommendation -- it is a consequence of the physics and the observed engineering pace. The trajectory projections assume the halving rate continues unchanged and no material wall is encountered earlier. SCALING CAVEAT: for SC platforms past T1*, the trajectory assumes the halving rate resumes from the current baseline -- this is only valid if qubit count and connectivity do not increase on the same substrate. See Substrate Inversion Amplifier.

Qubit Count Scaling Law

How gate error grows with qubit count -- derived from Landauer additivity, April 2026

A published gate error rate is measured on a single gate pair with neighboring qubits idle. Real circuits run gates across the full chip in parallel. The IAMPerformance scaling law answers: how does effective gate error change with qubit count and connectivity? From the Landauer additivity principle -- each qubit is an independent decoherence channel -- effective error scales with both qubit count and connectivity density. The crosstalk coupling is architecture-specific and substrate-state-dependent. Three regimes confirmed from published data: SC heavy-hex pre-T1*: scaling essentially free -- IBM Falcon to Eagle, 4.7x qubits, ~0% degradation. SC CZ post-T1*: scaling toxic -- Heron R2 to Nighthawk, +9% connectivity, +7.6% regression. QCCD ion trap: scaling negligible -- H2-1 to Helios, +75% qubits, -57% gate error. Rydberg neutral atom: essentially free -- Gemini 260 to 3000Q, 12x qubits, -63% error. The connectivity function is architecture-specific: SC grid grows with N; QCCD grows logarithmically; Rydberg stays flat. Derived April 3, 2026. Patents 64/012,720 and 64/014,568.

Substrate Inversion Amplifier

The crosstalk multiplier that activates when T1 crosses T1 -- derived April 2026*

The three-component decomposition shows that the crosstalk cost per added coupler is not constant -- it depends on the substrate's TLS saturation state relative to T1*. Pre-inversion (T1 < T1*): TLS defects are not saturated. Each new coupler couples through a clean junction environment. Scaling is essentially free. Post-inversion (T1 > T1*): TLS defects at every junction are saturated. Each new coupler now propagates crosstalk through a lossy medium. Scaling becomes substantially more costly. The Substrate Inversion Amplifier is the ratio of these two regimes -- a large multiplier that activates exactly when T1 crosses T1*. The IBM trajectory confirms it across nine published platforms: Falcon to Eagle (4.7x qubits, pre-T1*): ~0% degradation. Heron R2 to Nighthawk (+9% connectivity, post-T1*): +7.6% regression. Escape: Tantalum raises T1* to ~1300 us, returning the platform to the free-scaling regime. Derived April 3, 2026. Patents 64/012,720 and 64/014,568.

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-1 36Q has reached the Al/AlOx design ceiling. The next Rigetti generation requires a substrate material change.

Basis: Cepheus-1 36Q at 1.34x Al/AlOx floor. Cepheus-1 108Q at identical gate error (3.35x floor) confirms: scaling on Al/AlOx does not improve performance. Two independent data points confirm the ceiling.

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 to IonQ EQC trajectory). Wall projected early 2027.

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

Microsoft Majorana: if A > 0.002 upon first two-qubit gate error publication, topological protection is not outperforming standard SC transmon class. A < 0.001 would indicate topology is contributing meaningfully.

Basis: Binary test. Majorana n predicted ~1.19 (SC ECR class prior). Jul 2025 parity data: Z-loop readout error ~0.5%, X-loop ~16% (single-qubit parity, not gate). Two-qubit braiding not yet demonstrated. Nature editorial: paper does not represent evidence for MZMs. Prior 2018 Majorana paper retracted 2021. Framework will update immediately upon first published two-qubit gate error.

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.

IAM-2026-P012 Apr 3, 2026 **PENDING**

Any IBM SC system on Nb/AlOx with T1 > T1* = 325 us that increases connectivity density will show gate error regression. The next IBM generation beyond Nighthawk (N > 218 qubits) on the same substrate is predicted to reach A > 2.3x10^-3 -- worse, not better, than Nighthawk. The Substrate Inversion Amplifier means each added coupler now costs substantially more in crosstalk than it did in the Falcon-to-Eagle era. Escape route: substrate change to Tantalum (T1* = 1300 us) returns the platform to free scaling.

Basis: IAMPerformance qubit count scaling law. Derived April 3, 2026. Calibrated from Heron R2 -> Nighthawk transition: +9% connectivity increase caused +7.6% gate error regression despite better T1. IBM Falcon->Eagle on same substrate pre-T1*: 4.7x more qubits, ~0% degradation. The Substrate Inversion Amplifier activates when T1 crosses T1* -- crosstalk cost per coupler increases substantially. Substrate change to Tantalum (T1* = 1300 us) returns the platform to the pre-inversion free-scaling regime. Patents 64/012,720 and 64/014,568.

IAM-2026-P013 Apr 3, 2026 **PENDING**

Quantinuum Sol (192Q, 2027) will achieve A < 7.9x10^-4 -- equaling or improving on Helios 98Q. Specific prediction: A_Sol ~ 5.8x10^-4. QCCD scaling law: connectivity grows as log(192)/log(98) = 1.15x. Scaling degradation from qubit count alone: < 0.3% (alpha_QCCD ~ 0.02). Engineering improvement at the observed 2.2-year Ba+ halving rate dominates. Quantinuum Sol scales for free -- no substrate, no TLS ceiling, no Substrate Inversion. This prediction can also be stated as: Sol will NOT regress from Helios despite 96% more qubits.

Basis: IAMPerformance QCCD scaling law. alpha_QCCD ~ 0.02 calibrated from H2-1 (56Q) -> Helios (98Q): +75% qubits, -57% gate error. QCCD connectivity scales as log(N), g = 1 always (vacuum substrate, no TLS mechanism). Ba+ halving rate 2.2 years from Quantinuum arXiv:2511.05465. Filed April 3, 2026.

IAM-2026-P014 Apr 3, 2026 **PENDING**

IBM Tantalum substrate processor (when published): will scale like the Falcon-to-Eagle era, not the Heron-to-Nighthawk era. Prediction: at N > 200 qubits on Tantalum, gate error will NOT regress with connectivity increase. T1* for Tantalum = 0.65 x 2000 us = 1300 us -- IBM will need to reach T1 > 1300 us before the Substrate Inversion Amplifier engages. At IBM's current T1 improvement pace, that is approximately 4-6 chip generations away. If IBM Tantalum publishes a regression comparable to Nighthawk's (+7.6% for +9% connectivity) while T1 < 1300 us, this prediction is falsified.

Basis: IAMPerformance substrate-specific scaling prediction. Tantalum T1_free = 2000 us (Place et al. 2021, Nature Commun. 12:1779). T1* = 0.65 x 2000 = 1300 us. Prediction follows from the three-component model: pre-T1* regime, scaling free. Filed April 3, 2026. Filed April 3, 2026.

DATA INDEX

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

Platform	p(2Q)	A Score	T_op	Primary Source
Oxford Ionics EQC prototype	8.400e-05	8.400e-05	1 mK	arXiv:2510.17286, Hughes et al. (Oxford Ionics/IonQ), Oct 2025
Quantinuum Helios 98Q	7.900e-04	7.903e-04	1 mK	arXiv:2511.05465, Ransford et al. (Quantinuum), Nov 2025
Quantinuum H1-1	8.600e-04	8.604e-04	1 mK	Quantinuum blog, Apr 16 2024: 99.914(3)% fidelity
Google Willow 105Q	1.500e-03	1.501e-03	15 mK	Google Quantum AI, Nature 638:920 (Dec 2024)
Quantinuum H2-1 56Q	1.840e-03	1.842e-03	1 mK	Quantinuum quantinuum.com: 99.816(5)% fidelity
IBM Heron R2 156Q	2.000e-03	2.002e-03	15 mK	IBM QDC 2024; AbuGhanem (2024) doi:10.1007/s11227-025-07047-7
IBM Nighthawk 120Q	2.154e-03	2.156e-03	15 mK	IBM Quantum Platform announcement Jan 5 2026
IonQ Forte 36Q	4.000e-03	4.008e-03	1 mK	IonQ Forte Enterprise spec ionq.com/quantum-systems/forte-enterprise
Rigetti Cepheus-1 36Q	4.000e-03	4.008e-03	15 mK	Rigetti Q4 2025 earnings Mar 4 2026: 99.6% median CZ fidelity (36Q system)
QuEra Gemini 260Q	8.000e-03	8.032e-03	room temp	QuEra quera.com/gemini: >99.2% two-qubit fidelity
Rigetti Cepheus-1 108Q	1.000e-02	1.005e-02	15 mK	Rigetti Q4 2025 earnings Mar 4 2026: 99.0% median CZ fidelity (108Q chiplet system)
Rigetti Ankaa-3 84Q	1.000e-02	1.005e-02	15 mK	Rigetti GlobeNewswire Dec 23 2024: 99.0% median iSWAP fidelity
Microsoft Majorana 1	PENDING	PENDING	--	Aghaee et al., Nature 2025. doi:10.1038/s41586-025-08472-9

Platform	p(2Q)	A Score	T_op	Primary Source
Google Neutral Atom (2026+)	PENDING	PENDING	--	Google Quantum AI announcement March 24 2026

QAPE

Quantum Analytical & Performance Engine

IAMPerformance | Quantum Computing Performance Intelligence

IAMPerformance is an independent research initiative developing physics-based performance intelligence for information processing hardware.

QAPE is our quantum computing domain-specific commercial implementation of the IAM framework -- calibrated to the dominant noise mechanisms and material science of quantum computing hardware. The same underlying physics as the cosmological research. The same first-principles derivation. Different substrate.

Quantum computing hardware -- gate fidelity, error correction threshold, material wall analysis, architecture classification, competitive trajectory across all published platforms.

INPUTS

+ Platform name
+ Architecture class (8 classes)
+ Material substrate (13 including predicted, with physics notes)
+ Published p(2Q) gate error rate
+ Operating temperature (mK)
+ T1 coherence time (μs)
+ T2 coherence time (μs)
+ Qubit count
+ Year published
+ Publication date -- full YYYY-MM-DD for P-number stamping
+ Source / citation
+ Prior generation platform + p(2Q)
+ Prior generation year

LEVERS

+ Target temperature
+ Target substrate (13 options, live physics notes per material)
+ Participation ratio engineering factor
+ Target halving rate
+ Combined -- any combination simultaneously

OUTPUTS

+ A score = -ln(1-p)
+ Architecture parameter n
+ Material A-floor (IAMPerformance-derived)
+ Wall ratio -- current A / A-floor
+ Wall concern -- STALL / HIGH / MEDIUM / LOW
+ QEC threshold status + years to crossing
+ QEC crossing year
+ FT threshold status + years to crossing
+ FT crossing year
+ Full temperature sweep table (p at any T)
+ Cooling sensitivity rating and ratio
+ Temperature at which QEC is crossed
+ Temperature at which FT is crossed
+ Halving rate (observed or operator-entered)
+ 16-year A score projection table
+ Material floor reached year
+ Then/Now delta with regression detection
+ New A-floor after substrate change
+ New wall ratio and headroom gained
+ New QEC/FT timeline after acceleration
+ PR-adjusted effective floor
+ Combined lever scenario
+ T2/T1 anomaly detection and flag
+ Gate error decomposition: p_coh + p_mat + p_ctrl (SC architectures)
+ Substrate Inversion T1* status: PAST / APPROACHING / PRE-TRANSITION (SC only)
+ Motional Heating Inversion status: MHI timing for Ion A architecture
+ Architecture Dennard transition class and estimated crossing year
+ Validation: Ion A / Ion B / Neutral Atom confirm NOT substrate-limited
+ Qubit count scaling regime: pre-T1* (free) / post-T1* (toxic) / QCCD (negligible)
+ Substrate Inversion Amplifier g: alpha ratio post/pre-T1* (~16x when past T1*)
+ Scaling-adjusted trajectory flag: halving rate caveat for SC platforms past T1*
+ Optimal substrate prediction: which material change restores free scaling
+ Auto-generated P-number predictions (up to 5)
+ Tour guide commentary on every output block
+ Strategic punchline

VISUALIZATION

+ Threshold Proximity -- all platforms ranked by distance to QEC and FT, color-coded by wall concern
+ Temperature Sweep -- A(T) curve for one or two platforms, threshold crossings computed
+ Substrate Runway -- all platforms vs material floor, headroom visible at a glance

ANALYSIS

+ Benchmark Translation -- any published metric to IAMPerformance A score
+ Reverse Solver -- what gate error do you need to lead a competitor by a target year?
+ Portfolio Analysis -- optimal architecture hedge for speed now vs ceiling protection later
+ Acquisition Analysis -- does this target complement or compete with your platform?
+ QAPE Diagnostics -- full engineering diagnostic with auto-generated dated predictions

IAM

IAM QAPE (Quantum Analytical and Performance Engine)

Informational Actualization Model · v2.1 · Mahaffey (2026)

EXIT

PLATFORM SELECTION

Select a known platform

1 · IBM Nighthawk 120Q

Select a platform to load published data. Active platforms are pre-filled and ready to run.

Published gate error (pre-filled — override with your own data if needed)

0.00215

Published value pre-filled. Click Run Analysis to use this, or clear and enter your own value.

▾ OPTIONAL DATA (IMPROVES ANALYSIS)

T1 — qubit lifetime (microseconds)

e.g. 68

How long a qubit holds its state. Find it: look for T1 or relaxation time in the paper.

T2 — phase stability (microseconds)

e.g. 89

Usually shorter than T1. Find it: look for T2 or dephasing time.

Prior chip result (for improvement rate)

e.g. 0.00757, 15, 2024

Gate error, temperature (mK), year — of an older chip from the same company.

Improvement IAM Improvement Rate (years)

e.g. 2.0

IBM: 2.0 · Google: 3.3 · Quantinuum: 6.5 · Oxford: 0.1. Leave blank to use class default.

RUN ANALYSIS

IBM Nighthawk 120Q

SC Transmon — ECR gates (IBM, Rigetti) · 15 mK · p_2Q = 2.150e-3 · The Quantum Insider, Jan 13 2026; AbuGhanem (2024)

IAM PERFORMANCE METRIC

0.002152

Lower is better. World best 2Q gate: 8.40e-05 (Oxford I 2025)

IAM COMPETITIVE INDEX

1	8.400e-5	Oxford Ionics EQC (20
2	0.000790	Quantinuum Helios 98Q
3	0.001501	Google Willow 105Q (2
4	0.002002	Quantinuum H2 56Q (20
5	0.002152	**IBM Nighthawk 120Q**
6	0.003005	IonQ Forte 36Q (2024)
7	0.003005	QuEra Gemini 3000Q (2
8	0.007025	Rigetti Ankaa-3 84Q (

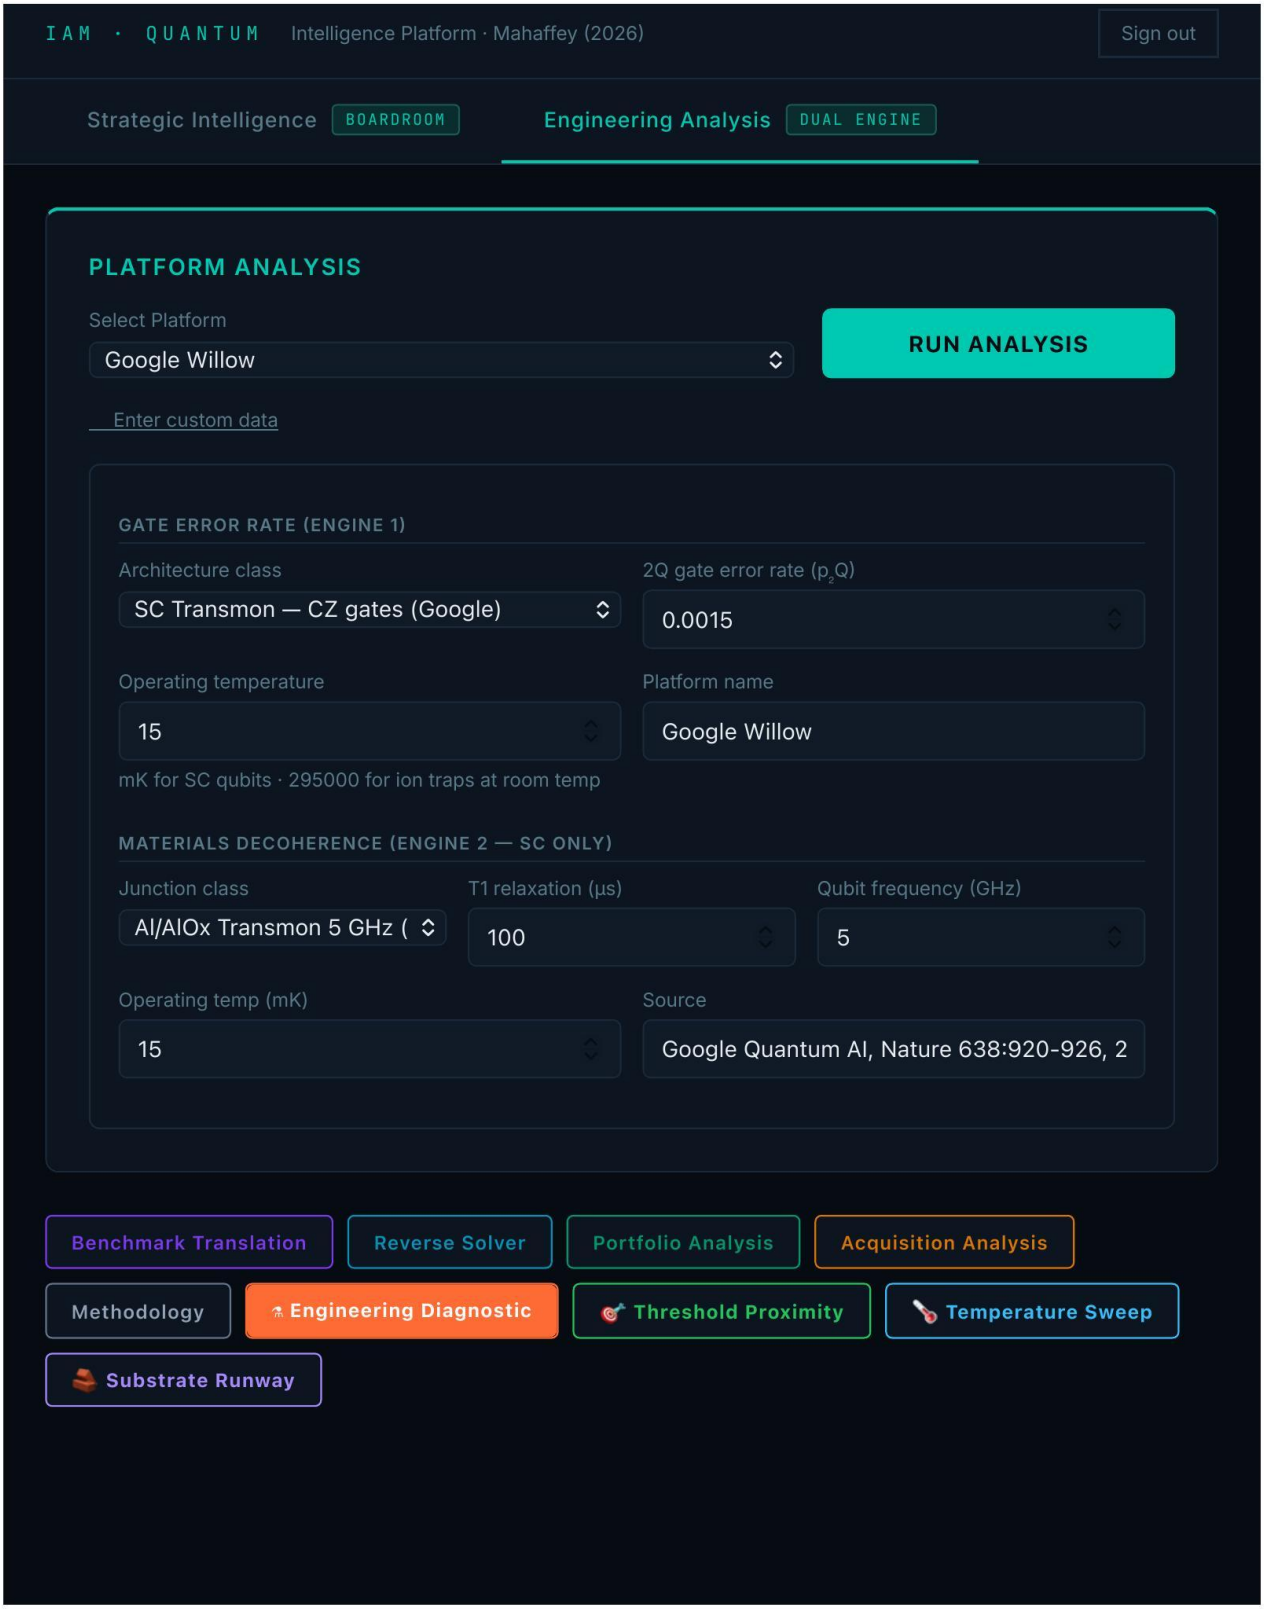
IAM · QUANTUM
Intelligence Platform · Mahaffey (2026)
Sign out
Strategic Intelligence
BOARDROOM
Engineering Analysis
DUAL ENGINE
PLATFORM ANALYSIS
Select Platform
Google Willow
RUN ANALYSIS
Enter custom data
GATE ERROR RATE (ENGINE 1)
Architecture class
SC Transmon — CZ gates (Google)
2Q gate error rate (p2Q)
0.0015
Operating temperature
15
Platform name
Google Willow
mK for SC qubits · 295000 for ion traps at room temp
MATERIALS DECOHERENCE (ENGINE 2 — SC ONLY)
Junction class
Al/AlOx Transmon 5 GHz (
T1 relaxation (µs)
100
Qubit frequency (GHz)
5
Operating temp (mK)
15
Source
Google Quantum AI, Nature 638:920-926, 2
Benchmark Translation
Reverse Solver
Portfolio Analysis
Acquisition Analysis
Methodology
Engineering Diagnostic
Threshold Proximity
Temperature Sweep
Substrate Runway

THE OUTPUTS

Submit a run and the engine returns everything it can derive -- every metric laid out as a strip with plain-language explanations below each value. The Then/Now regression detector fires automatically if prior generation data was entered. Every value is specific, dated, and falsifiable. Every output block carries a tour guide paragraph explaining what the numbers mean and why they matter. The strategic punchline closes every run.

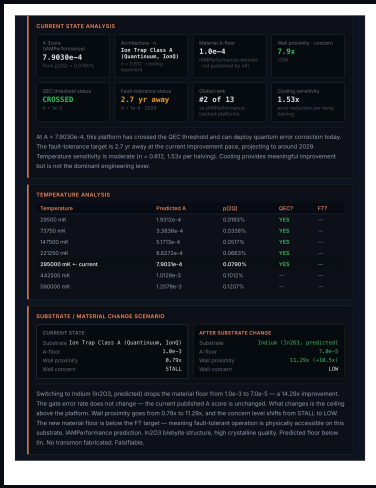

QAPE output for Quantinuum Helios 98Q: A = 7.903e-4, n = 0.612 (Ion A class), Material A-floor = 1.0e-4 (Ba+ ion trap, IAMPerformance-derived), Wall proximity = 7.9x (LOW concern -- substantial runway remaining), QEC threshold: CROSSED, FT target: ~6.6 yr (projected 2032), Global rank: #2 of 13 tracked platforms, Cooling sensitivity: 1.53x per halving.

SUBSTRATE RUNWAY

Every active platform positioned relative to its material A-floor simultaneously. The gap between the current A score and the floor is the engineering runway -- how much improvement is physically possible without changing the substrate. Color-coded by wall concern. All A-floor values are QAPE-derived from first principles. Not published by manufacturers. Not available in any public database.

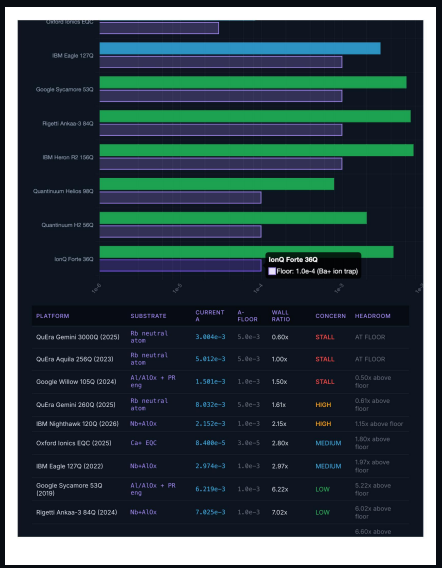

Substrate Runway: bars show current A score, material floor line shows physical ceiling. Green = LOW concern (headroom available), amber = HIGH (wall approaching), red = STALL (at or past the material ceiling). IonQ EQC and Helios: most runway. Willow and Cepheus: wall-imminent.

THRESHOLD PROXIMITY AND TEMPERATURE SWEEP

Threshold Proximity answers the question every serious buyer is actually asking: which platform is closest to being genuinely useful for protected computation? Every tracked platform is ranked by its distance to the QEC and FT thresholds, color-coded by wall concern, with years to each threshold and architecture class shown. Temperature Sweep plots the A(T) curve for one or two platforms across any temperature range -- the QEC and FT crossing temperatures are computed automatically. The thermal prediction made visible.

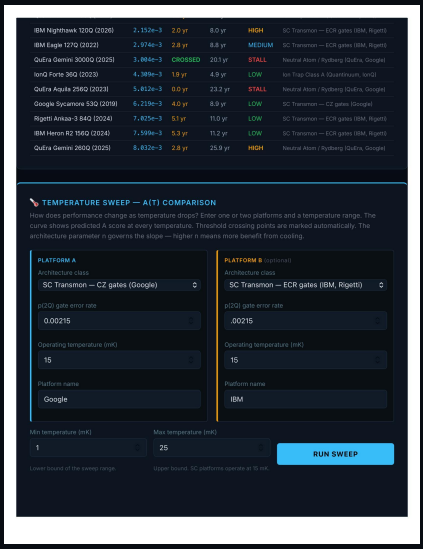

Left: Threshold Proximity -- 13 platforms ranked by A score, years to QEC, years to FT, wall concern, architecture class. Right: Temperature Sweep -- compare two platforms on the same A(T) curve over any operator-defined temperature range. Threshold crossing temperatures computed.

REVERSE SOLVER AND PORTFOLIO ANALYSIS

The Reverse Solver answers the competitive gap question: what gate error rate do you need to lead a specific competitor by a target year, and how long does it take at your current pace? It then ranks the available engineering levers by impact and implementation cycle time. The Portfolio Analysis identifies the optimal architecture hedge -- which combination of architecture classes gives the best balance of near-term performance advantage and long-term ceiling protection.

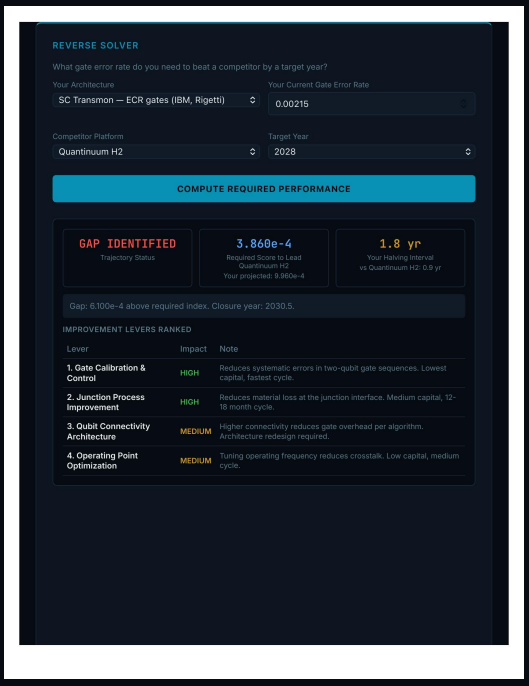

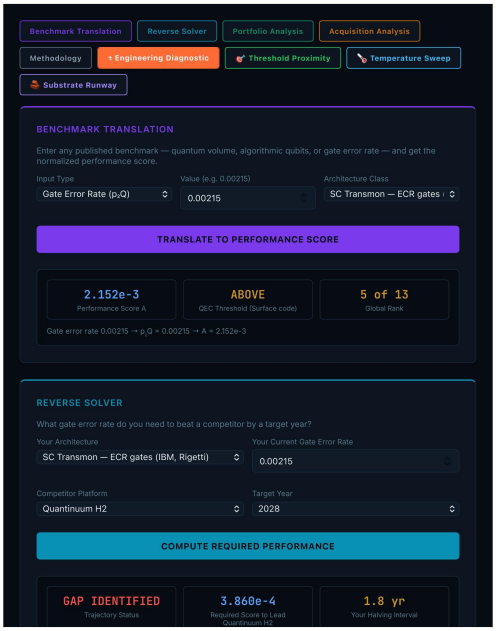

Left: Reverse Solver -- gap identified, required A score shown, engineering levers ranked by impact and implementation cycle time. Right: Portfolio Analysis -- optimal hedge result showing the strongest strategic position combines an above-threshold architecture with an open-runway architecture.

Enter any data you have. Leave fields blank if unpublished — the framework derives everything it can from what you provide. Engineering levers let you model substrate changes, temperature scenarios, acceleration programs, or any combination before spending a dollar in the lab. Every output is a dated, falsifiable IAMPerformance prediction.

SECTION 1 — PLATFORM IDENTITY

Who are we analyzing? Select from the library or enter a custom platform. Architecture class is the single most important field — it determines the temperature sensitivity parameter n and the applicable QEC threshold. If you're not sure, SC CZ covers Google and IBM's latest generation; Ion A covers Quantinuum and IonQ Forte.

Platform (select or enter custom below)

Quantinuum Helios 98Q

Selecting a platform pre-fills all known published values. Any field can be overridden.

Platform name

Quantinuum Helios 98Q

Used in output labels and PDF export. Can be anything.

Architecture class *

Ion Trap Class A (Quantinuum, IonQ)

SC CZ — Google Willow, IBM Heron/Nighthawk, Rigetti Cepheus (n=1.449, highest thermal sensitivity) | **SC ECR** — IBM Eagle, Rigetti Ankaa-3 (n=1.189) | **Ion A** — Quantinuum, IonQ Forte (n=0.612, low thermal sensitivity) | **Ion B** — Oxford Ionics EQC (n=1.200) | **Neutral** — QuEra Gemini (n=0.351, cooling irrelevant) | **Topological** — Microsoft Majorana (n=1.189 predicted)

Year published

2026

Used for trajectory year calculations (QEC crossing year, FT crossing year).

SECTION 2 — PUBLISHED GATE DATA

Enter what the platform has published. The gate error rate p(2Q) is the only required field. Everything else improves the analysis. If a value is not published, leave it blank — the framework will note what it cannot derive.

Two-qubit gate error rate p(2Q) *

0.00079

The probability a two-qubit gate produces the wrong state. This is the single number the IAMPerformance A score is derived from. For IBM platforms, use EPLG@100q. For neutral atoms, use the published two-qubit Rydberg gate fidelity error.

Operating temperature (mK)

295000

SC qubits: typically 15 mK (dilution refrigerator base temperature). Ion traps: 1 mK for the ion chain. Neutral atoms: room temperature (enter 295000). This is the reference point for all thermal predictions.

Improvement rate — halving years

e.g. 2.0

How many years for the A score to halve at the current program pace. SC CZ class typical: 2.0-3.3 yr. Ion A class typical: 2.2 yr (Ba+), 7.9 yr (Yb+). EQC: 0.63 yr. If unknown, the framework uses the class default.

T1 — relaxation time (μs) optional

e.g. 350

How long a qubit holds its energy state. IBM Nighthawk: 350 μs (highest in fleet). Google Willow: 68 μs. If entered alongside T2, the framework checks for the T2/T1 anomaly — T2 should not exceed T1 for CZ-gate SC platforms.

T2 — phase coherence time (μs) optional

e.g. 89

How long a qubit maintains phase coherence. Standard physics: T2 ≤ 2×T1. Google Willow reports T2 = 89 μs with T1 = 68 μs — T2/T1 = 1.309, flagged as anomalous. The diagnostic engine detects and reports this automatically.

Qubit count optional

e.g. 120

Number of physical qubits on the chip. Not used in the A score calculation — gate

Publication date for P-number dating

e.g. 2026-04-03

The date on which this data was published. Used to timestamp the auto-

Source / citation optional

e.g. arXiv:2510.17286

Primary source for the published gate error rate. Recorded in the predictions

A FINAL NOTE

Quantum computing is the most consequential engineering race in modern history. The teams working on it are doing genuinely hard things -- building systems that operate at temperatures colder than deep space, chasing error rates that would have seemed fictional a decade ago, making multi-hundred-million-dollar bets on physics that is still being written. Every engineer on every one of these programs deserves better than guesswork about where their platform stands and what the physics actually allows.

What we have been fortunate enough to discover is a set of physical relationships that answer those questions from first principles. Not from fitted curves. Not from historical analogies. From the underlying physics of what it actually means for a qubit to lose coherence -- and what it means to get it back. Those relationships let us say, with confidence, which architecture classes have room to grow and which have reached their material ceiling. Which engineering lever actually moves performance for a given hardware type. Which platforms are months from fault tolerance and which are a decade away. All from a single published number, in seconds.

The depth of that derivation matters and deserves to be stated plainly. The framework is grounded in the minimum energy cost of any irreversible information event -- the same physical law that governs the cosmic horizon, the transistor junction, and the qubit gate through one expression, not by analogy. The architecture classification that places every platform in a class comes from independent experiments conducted on molecules at thousands of Kelvin in 2004 -- experiments that had nothing to do with quantum computing -- whose thermal exponents transferred across eight orders of magnitude in temperature and placed superconducting qubits correctly before those qubits existed in their current form. The material floor for each substrate comes from oxide stoichiometry and defect site physics, not from fitting a line through hardware data. And the decomposition that separates a published gate error into its three physically independent components -- coherence, substrate, and control overhead -- means that for the first time, an engineer can look at a single published number and know exactly which part of it is fixable, which part requires a new material, and which part reflects the irreducible floor of the gate mechanism itself. That was an opaque number before. It is not anymore.

The value is not the ranking table. The value is the months of experimental campaigns that never need to happen because the answer was already in the data. The cooling investment that gets redirected before the dilution refrigerator arrives, because the framework already showed it wouldn't cross the threshold. The substrate decision made in a day instead of a year, because the physical ceiling was derivable before a single wafer was spun. The qubit count decision informed before a single wafer is taped out, because the framework can now tell you whether adding more qubits on your current substrate will improve performance or regress it -- and exactly which material change restores the runway. Every one of those recovered months is a month spent actually advancing toward error-corrected computation instead of rediscovering a constraint the physics already knew.

IAMPerformance is an independent research initiative. There is no institutional affiliation, no investor agenda, and no stake in any of the platforms tracked in this publication. The sole mission, over the coming weeks and months, is to continue refining and optimizing QAPE -- expanding the architecture classes it covers, sharpening its predictions as new data publishes, and building it into a tool that the platforms themselves can use to accelerate the journey. Every team working toward fault-tolerant quantum computing deserves a navigation instrument that tells them where they actually are, not where they hoped to be. This publication is one step toward that. The goal was never to rank hardware. The goal is to get quantum computing working -- correctly, efficiently, and for the benefit of the people who will eventually use it for things none of us can fully imagine yet. We are fortunate the physics cooperated.