Certified randomness from quantum supremacy (Nature)

Portability dossier - app

agQSL portability pipeline

2026-06-07

Field	Value
Slug	`2026-nature-certified-randomness-quantum-supremacy-nature`
Source	journal
Link	https://www.nature.com/articles/s41586-025-08737-1
Category	app
Triaged	2026-06-07 by port_until_julien_parallel
Bootstrapped	2026-06-07

Paper: fetch failed (see `paper.url`)

Contents

1 Source

1.1 Domain classification

`cryptography` (certified randomness / randomness expansion). Secondary: sampling-based quantum-advantage benchmarking (random circuit sampling), which supplies the computational-hardness primitive but is not itself the application.

Identity note for Julien. Despite the slug, the Nature DOI `s41586-025-08737-1` resolves to "Certified randomness using a trapped-ion quantum processor" (Liu, Shaydulin, Niroula, DeCross, Hung et al., 2025), the experimental realisation on Quantinuum's H2-1, not Google's superconducting supremacy run. Fingerprint trapped-ion hardware, not superconducting.

1.2 Expert persona for Julien

A quantum cryptographer or quantum-information theorist fluent in random circuit sampling and linear cross-entropy benchmarking (LXEB/XEB), the Aaronson-Hung certified-randomness protocol, min-entropy certification, and randomness extractors. Should have calibrated intuition for trapped-ion device metrics (two-qubit gate fidelity, all-to-all QCCD connectivity, mid-circuit measurement, ion-shuttling overhead) and, crucially, for the classical co-processing side: this fingerprint has two halves, the 56-qubit H2-1 sampler and the leadership-class supercomputers ($\sim 1.1 \times 10^{18}$ FLOP/s sustained) that run verification. Someone comfortable reading fidelity and circuit-geometry numbers out of a Nature supplementary file.

1.3 Related prior work (brief bibliography)

- Aaronson & Hung, "Certified Randomness from Quantum Supremacy", arXiv:2303.01625, STOC 2023, doi:10.1145/3564246.3585145. The theory protocol this paper realises.
- Arute et al., "Quantum supremacy using a programmable superconducting processor", Nature 2019, doi:10.1038/s41586-019-1666-5. The same RCS primitive on a different vendor (Google superconducting); the natural portability comparison.
- Moses et al., "A Race-Track Trapped-Ion Quantum Processor", Phys. Rev. X 13, 041052 (2023), arXiv:2305.03828. The H2 device platform.
- Aharonov, Gao, Landau, Liu & Vazirani, "A polynomial-time classical algorithm for noisy random circuit sampling", arXiv:2211.03999, STOC 2023. Bears on the classical spoofing cost that underwrites the security claim and the verification budget.
- Bierhorst et al., "Experimentally generated randomness certified by the impossibility of superluminal signals", Nature 556, 223 (2018), doi:10.1038/s41586-018-0019-0. An alternative (Bell-test, device-independent) route to certified randomness, useful for framing what this protocol does and does not assume.

1.4 Difficulty estimate

`high`. The Nature page is paywalled and the bootstrap fetch failed (HTML landing page, see `paper.url`), but the full preprint with supplementary detail is open at arXiv:2503.20498. The extraction is heavy regardless: the result couples an exotic cryptographic stack (LXEB certification, min-entropy bookkeeping, randomness extraction, restricted-adversary assumptions; 71,313 certified bits) to a two-sided hardware fingerprint (56-qubit H2-1 plus exascale classical verification), and the device-level numbers Julien needs will sit in the SI rather than the main text. Recommendation to the triager: drop the arXiv:2503.20498 PDF into the folder as `paper.pdf` before Julien runs, so device fidelities and circuit specifications are readable without the paywall.

2 Extraction

Identity note. The dossier slug and title read "Certified randomness from quantum supremacy (Nature)", but the Nature DOI `s41586-025-08737-1` resolves to *Certified randomness using a trapped-ion quantum processor* (Liu, Shaydulin, Niroula, DeCross, Hung et al., 2025). The fingerprint below is for the Quantinuum H2-1 trapped-ion realisation, not Google's superconducting supremacy run. Numbers are taken from the open-access preprint arXiv:2503.20498 (TeX source, main text plus Supplemental Material), which corresponds to the same content as the paywalled Nature article. See Caveats.

2.1 What the paper does (one paragraph)

The work generates random bits whose unpredictability can be certified by a classical client talking to an untrusted remote quantum computer over the internet, a task provably impossible for a purely classical device. The client pseudorandomly generates 56-qubit "challenge" random circuits from a short secret seed, sends them in batches to the Quantinuum H2-1 trapped-ion processor, and demands each batch back within a fixed time. It then scores a small subset of returned bitstrings using linear cross-entropy benchmarking (XEB). A high XEB score combined with a response time too short for any realistic classical simulator certifies that the server must have used a quantum computer, and hence that the returned bits contain genuine entropy. The classical hardness rests on random circuit sampling: verifying one challenge circuit by exact tensor-network contraction took about 100 seconds on the full Frontier supercomputer. Running verification across four supercomputers at a combined 1.1×10^{18} FLOP/s, the authors certify 71,313 bits of smooth min-entropy against a restricted near-term adversary, then extract 71,273 near-uniform bits, achieving certified randomness expansion from a 32-bit seed.

2.2 Quantum hardware used

- **Vendor / machine:** Quantinuum H2-1 trapped-ion quantum processor, accessed remotely over the internet.
- **Qubit count:** 56 used / 56 available. All $n = 56$ qubits are algorithmic; there is no error-correction encoding, so logical and physical counts coincide (Table I).
- **Connectivity:** All-to-all, realised by the QCCD (quantum charge-coupled device) architecture through physical ion shuttling. The two-qubit gate pairings in each layer come from an edge colouring of a random d-regular graph on n nodes, so arbitrary qubit pairings are required (SI Sec. IV C).
- **Gate set:** Native two-qubit gate $U_{\mathrm{ZZ}}(\pi/2)$ (arbitrary-angle ZZ, the native Quantinuum entangler), interleaved with layers of arbitrary single-qubit $SU(2)$ rotations on all qubits (SI Sec. IV C).
- **Notable features leveraged:** Mid-circuit measurement and reset (two circuits are "stitched" into one job, joined by a layer of mid-circuit measurement and reset, to amortise latency); QCCD ion shuttling for reconfigurable connectivity; very low per-circuit overhead, which makes single-shot-per-circuit execution viable without a time penalty (main text; Methods Sec. Protocol Details).

2.3 Computational primitive

Sampling. Specifically random circuit sampling (RCS): the quantum processor draws one bitstring per random circuit from that circuit's output distribution. The application built on top is certified randomness / randomness expansion, but the quantum primitive exercised on hardware is RCS. The protocol is non-variational: there is no classical parameter-optimisation loop.

2.4 Resource fingerprint

Metric	Value	Source (page / equation / SI)
Qubits (logical)	not applicable (no error correction; $n = 56$ algorithmic qubits)	Table I; main text
Qubits (physical)	56	Table I
Circuit depth	$d = 10$ entangling layers (21 gate layers total: 11 $SU(2)$ layers interleaved with 10 U_{ZZ} layers)	SI Sec. IV C (line 1098)
2Q-gate count (total)	280 per circuit ($10 \times 56/2$); 616 single-qubit gates per circuit (11×56)	SI Sec. IV C (line 1098)
Measurement shots	1 sample per circuit (single-shot); $M = 30{,}010$ accepted samples from 60,952 submitted circuits (1,993 batches, 984 successful)	Table I; main text; SI Sec. IV D
Classical loop iter.	not applicable (non-variational RCS; no optimisation loop)	n/a

Metric	Value	Source (page / equation / SI)
Wall-clock runtime	Quantum: cumulative device time 64,652 s ($\approx$ 18 h) for accepted samples; $t_{\mathrm{QC}} = 2.154$ s per sample ($\approx$ 1 bit/s certified). Classical verification: 100.3 s per circuit on full Frontier, 153.5 s per circuit on full Summit; $\approx$ 370,000 effective Frontier node-hours ($\approx 2.8 \times 10^{23}$ FLOPs) over four supercomputers for $m = 1{,}522$ verified circuits	Table I; main text; SI Sec. IV A, IV E

Additional certified-randomness and hardness numbers (all from main text and Table I unless noted):

- Per-circuit exact simulation cost $\mathcal{B} = 90 \times 10^{18}$ FLOPs on Frontier (SI Sec. IV A rounds this to $\approx 10^{20}$ FLOPs per circuit at $\approx 50\%$ numerical efficiency).
- Adversary classical power $\mathcal{A} = 0.897 \times 10^{18}$ FLOP/s (one Frontier); certified result quoted against an adversary 4$\times$ Frontier.
- Combined sustained verification performance 1.1×10^{18} FLOP/s (Frontier 897 PFLOPS at 45% efficiency plus Summit 228 PFLOPS at 59%; Perlmutter and Polaris also used).
- Thresholds: XEB threshold $\chi = 0.3$, time-per-sample threshold $t_{\mathrm{threshold}} = 2.2$ s, batch cutoff $T_{b,\mathrm{cutoff}} = 2.5 \times 2b$ s, batch sizes $b \in \{15, 20\}$.
- Measured $\mathrm{XEB}_{\mathrm{test}} = 0.32$; estimated genuine-server circuit fidelity $\phi \gtrsim 0.3$ on depth-10 circuits.
- Certified smooth min-entropy $H_{\mathrm{min}}^{\varepsilon_{\mathrm{s}}} = 71{,}313$ bits at $\varepsilon_{\mathrm{sou}} = 10^{-6}$ ($\varepsilon_{\mathrm{s}} = \varepsilon_{\mathrm{sou}}/4$), with $Q_{\mathrm{min}} = 1{,}297$ certified quantum rounds; 71,273 bits extracted via a Toeplitz (Cryptomite) seeded extractor; 32-bit seed implies net randomness expansion.

Verbatim circuit definition (SI Sec. IV C): an n-qubit depth-d circuit $C_{n,d}$ has d entangling layers, each a random set of $n/2$ disjoint $U_{\mathrm{ZZ}}(\pi/2)$ gates, sandwiched by layers of random $SU(2)$ gates on all n qubits, with the entangling arrangement set by edge colouring of a d-regular graph. The XEB score (Eq. 1) is

$$\mathrm{XEB}_{\mathrm{test}} = \frac{2^n}{m} \sum_{i \in \mathcal{V}} p_{C_i}(x_i) - 1, \qquad p_C(x) = |\langle x|C|0\rangle|^2.$$

2.5 Assumptions and results

- **Claimed.** Using H2-1 plus exascale classical verification, the authors experimentally realise an RCS-based certified-randomness protocol and certify 71,313 bits of smooth min-entropy (extracting 71,273 near-uniform bits) against a restricted near-term adversary, demonstrating a beyond-classical application of a gate-based digital quantum computer.
- **Error bars / noise assumptions.** Security is conditional on a restricted adversary model (SI Sec. III C): the server splits its M rounds into Q genuine quantum rounds and $M - Q$ classically simulated rounds, performs no postselection, never oversamples a circuit, holds a perfect-fidelity quantum computer, and has classical power bounded by $\mathcal{A}$ (taken as 4× Frontier). The central computational assumption is that no practical classical algorithm can spoof the XEB test for this circuit family. The genuine-server device fidelity $\phi \gtrsim 0.3$ is estimated from mirror benchmarking and component-gate fidelities reported in DeCross et al. 2024 (arXiv:2406.02501), not measured afresh here; at $\chi = 0.3$, $t_{\text{threshold}} = 2.2$ s the genuine-server failure probability is $p_{\text{fail}} \approx 50\%$. Soundness $\varepsilon_{\text{sou}} = 10^{-6}$.
- **Comparison to classical baseline.** The classical baseline is the security argument itself. The genuine server returns a sample in $t_{\text{QC}} = 2.154$ s, whereas exact tensor-network contraction of one challenge circuit costs $\mathcal{B} = 90 \times 10^{18}$ FLOPs, about 100.3 s on the full Frontier supercomputer (the world's most powerful at the time). The success condition is the fidelity gap $\phi \gg \mathcal{A} \cdot t_{\text{threshold}}/\mathcal{B}$ (Eq. 3): the circuit must be too costly for the adversary to simulate at high fidelity within $t_{\text{threshold}}$.

2.6 Portability flags

Features this paper leans on that bear on portability to other substrates:

- **All-to-all connectivity at 56 qubits is structural, not incidental.** The challenge-circuit family is an edge colouring of a random d-regular graph and was chosen (following DeCross et al. 2024) precisely because arbitrary, geometry-free qubit pairings push up classical simulation cost (high tensor-network treewidth) while staying shallow enough for high fidelity. On fixed-connectivity hardware (heavy-hex, linear, fixed 2D grids) realising these pairings needs SWAP routing, which inflates the 280 two-qubit gates per circuit substantially and degrades fidelity. The security margin depends on this; a routed version on a planar device is both easier to simulate classically and lower fidelity.
- **Native arbitrary-angle $U_{\text{ZZ}}(\pi/2)$ two-qubit gate.** Substrates whose native entangler is CZ or CNOT must decompose, and substrates without an arbitrary-angle ZZ primitive pay extra depth and error per entangling layer.
- **Mid-circuit measurement and reset.** Used to stitch two circuits per job to reduce per-job latency. Not required for protocol correctness, but the quoted throughput ($t_{\text{QC}} = 2.154$ s) depends on it; substrates without MCM+reset would see different timing, which feeds directly into the security threshold.
- **Very low per-circuit overhead enabling single-shot execution.** The protocol is deliberately strengthened to one sample per circuit because the QCCD device has minimal load/run overhead, so single-shot circuits are not penalised relative to many-shot sampling. On hardware where circuit loading dominates, the many-shot variant of Aaronson-Hung (entropy $\Omega(n)$ from many samples of one circuit) would be preferable, changing the protocol shape.

- **QCCD ion shuttling overhead is inside the timing budget.** The 2.154 s per sample includes shuttling, re-cooling, and network communication. Any substrate ports the timing budget, not just the gate count, because the security test is a wall-clock test ($t_{\text{threshold}}$), not a depth test.
- **The application is two-sided and needs exascale classical co-processing.** Certification required four DOE supercomputers at 1.1×10^{18} FLOP/s sustained for verification of $m = 1{,}522$ circuits. This half is independent of the quantum substrate but is essential to the application; any portability analysis of "certified randomness" must cost the classical verifier, not only the quantum sampler.

2.7 Caveats

- **Source.** The Nature article (`s41586-025-08737-1`) is paywalled and the dossier bootstrap fetch returned an HTML landing page (see `paper.url`); no `paper.pdf` was present in the folder. All numbers here come from the open-access preprint arXiv:2503.20498, downloaded as the arXiv TeX source bundle (main text `ms.tex`, supplement `supplement.tex`, macros `Macro_definitions.tex`). The published Nature version may differ in section numbering or minor wording; the quantitative results are expected to match.
- **Numbers not reported in this paper.** Per-gate single-qubit and two-qubit fidelities, T_1/T_2 coherence times, and gate durations are not given; the circuit-level fidelity $\phi \gtrsim 0.3$ is the only device fidelity quoted, and it is sourced to DeCross et al. 2024. These are marked accordingly and were not invented.
- **Minor internal inconsistency.** Total accepted device time is 64,652 s in Table I (and the source macros) but 64,641 s in the SI extraction section (line 1145). The Table I value is cited above.
- **Adversary scope.** The certified entropy holds only against the restricted adversary class and additional assumptions in SI Sec. III C; the authors explicitly state the bit rate, soundness, restricted model, and assumptions limit immediate production deployment.

3 Portability matrix