SCORECARD 005 · CLAIMS LEDGER

Rack-scale vendor survey

Six commercially available rack-mounted quantum systems, scored on Effective Computational Width against the classical simulation floor. Published August 2026.

“Rack-mounted quantum systems now available for on-premises purchase deliver computational capability beyond classical simulation.”

CLAIM CLASS	Market-interpretation claim. Graded as the generic formulation circulating in buyer conversations. Extends Scorecard 002 from two systems to the full commercially available field.
VERDICT	**FALSE — as stated.** No system surveyed reaches the classical simulation floor by any measure. Every entry sits below 30 effective qubits; the largest sits at approximately 15. None enters the contested 30–50 band. Deployability across the category is real, engineered, and in several cases genuinely novel.
SCOPE	Grades the market claim. Named systems appear as evidence base and their published specifications are assessed as published. Findings concern documentation and architecture, not conduct. No intent is asserted or implied.
REUSABLE TEST	Effective Computational Width under each architecture's binding constraint, compared to the classical simulation floor. Headline number is scored separately under the addressability test.

What the survey found

Six systems, five modalities, three continents. Every vendor has solved a real engineering problem, and several have solved one that was open five years ago: room-temperature operation on a wall outlet, cryogenics relocated inside a chassis, telecom-component photonics in a standard rack, a trapped-ion machine that ships to a customer site. Those are not marketing claims. They are delivered, deployed, and in several cases independently operated by national laboratories.

None of it changes what the machines compute. The classical simulation floor is a property of problem size, not of enclosure design, and the entire surveyed field sits below it by a wide margin. The widest single addressable register found in the survey is approximately 15 effective qubits, inferred from a published Quantum Volume figure. Exact state-vector simulation of a general 15-qubit circuit requires roughly 0.5 MB of memory.

A second, narrower finding runs through the survey and is recorded separately for each entry: the headline number published by a vendor frequently counts something other than qubits addressable in a single circuit — cores in a chassis, optical modes, photons under temporal multiplexing, or the dimension of the state space. Each instance is scored under the addressability test rather than characterised as a disclosure failure.

The classical floor used here

A system clears the floor when it executes, as a single addressable register, circuits whose width *and* entanglement structure jointly defeat exact state-vector simulation and tensor-network contraction on commodity hardware. Bands: at or below 30 qubits, unambiguously below the floor at any depth; 30 to 50, contested and requiring stated circuit structure; above 50 as one register at sufficient depth, above the floor. This supersedes the 40-qubit threshold published in Scorecard 002; see status history.

Evidence base, scored on the honest metric

ECW is published as a range with assumptions. Where a modality does not admit a width figure, the cell states why rather than forcing a number.

System	Modality	Headline	ECW (range)	Binding constraint	Floor
SAXON Q SXQ128	NV diamond, room temp	128 qubits AGGREGATE	**8** *single core; no demonstrated inter-core link*	Register composability	Below
Equal1 RacQ	Silicon spin, 0.3 K on-chip	6 qubits ADDRESSABLE	**4–6** *6 nearest-neighbour; 4 all-pairs at 99.3% 2Q*	SWAP-inflated depth vs coherence	Below
Quantum Brilliance QB-QDK2.0	NV diamond, room temp	Not stated UNDISCLOSED	**2–?** *2 peer-reviewed on prior generation; QDK2.0 unstated*	Register composability	Below
AQT IBEX Q1	Trapped ion Ca-40, room temp	12 qubits ADDRESSABLE	**7–12** *12 physical, all-to-all; 7 effective from QV 128*	Wall-clock throughput	Below
AQT LYNX	Trapped ion, room temp	QV 32768 ADDRESSABLE	**~15** *inferred from QV; qubit count not published*	Wall-clock throughput	Below
ORCA PT-1 / PT-2	Time-bin photonic, room temp	8 / ~90 qubits NON-COMPARABLE	**Not defined** *not a universal gate machine; width is not the operative variable*	Problem class	N/A
Quandela Belenos (MosaiQ-12)	Photonic gate model, QD sources	12 qubits / 24 modes SEE NOTE	**Undetermined** *12 nominal; success probability per gate not published*	Success-probability compounding	Below

Scoring note. Systems are not scored on identical criteria, because the modalities do not share a binding constraint and the disclosures are not identical. Where a figure is undisclosed it is marked undisclosed and no value is inferred. Every ECW figure carries the assumptions under which it was computed. The *Floor* column reports position relative to the 30-qubit lower band only; no system in this survey required the contested-band procedure.

Published fidelity and disclosure state

System	1Q fidelity	2Q fidelity	Readout	Evidence grade
SAXON Q SXQ128	99.92%	UNDISCLOSED	Not published	VENDOR
Equal1 RacQ	99.9% @ 140 ns	99.3% @ 200 ns, interleaved	99% @ 10 us, parity	VENDOR
Quantum Brilliance QB-QDK2.0	UNDISCLOSED	UNDISCLOSED	UNDISCLOSED	—
AQT IBEX Q1	99.966% *error (3.4 +/- 1.0)e-4*	98.7% average *error (1.3 +/- 0.3)%; min 0.9%, max 1.8%*	Not separately published	VENDOR + PEER
AQT LYNX	UNDISCLOSED	UNDISCLOSED	UNDISCLOSED	VENDOR *QV only*
ORCA PT-1 / PT-2	N/A — not gate model	N/A — not gate model	Detector efficiency not published	PEER *architecture*

System	1Q fidelity	2Q fidelity	Readout	Evidence grade
Quandela Belenos	UNDISCLOSED	UNDISCLOSED *gates are probabilistic; fidelity alone would not characterise*	UNDISCLOSED	VENDOR

Disclosure asymmetry is the survey's most consistent finding. Two of seven systems publish a two-qubit gate fidelity. Gate speed is published by two. No vendor in the survey publishes an end-to-end wall-clock figure for a named workload. Non-disclosure is recorded as a fact about the documentation and carries no inference about the underlying value.

Entry notes

SAXON Q SXQ128 — register composability

Sixteen cores of eight qubits, published as 128. Scored **AGGREGATE**: the headline sums resources that are not composable into one register on the published architecture. Parallel independent registers deliver throughput of 8-wide circuits, not one 128-wide circuit, and no scheduler closes that gap without entanglement between cores. The vendor position that cores communicate is recorded and converted to a kill condition rather than adjudicated here: classical measurement and feed-forward between cores does not produce a joint register, and the distinguishing evidence would be a composed circuit wider than one core, executed end to end, with its fidelity published. Facilities requirement genuinely eliminated — wall outlet, no cold room. Two-qubit fidelity undisclosed.

Equal1 RacQ — routing tax

Six qubits, honestly published as six. The headline passes the addressability test, which in this survey is the exception rather than the rule. The constraint is topology: on a linear six-chain, entangling the end qubits requires four SWAPs, which is twelve CNOTs before the intended gate fires. At 99.3% two-qubit fidelity, 0.993 raised to the thirteenth power is approximately 91.2%, and thirteen 200 ns gates is roughly 2.6 microseconds of circuit time against the coherence budget — two distinct costs, reported separately. Cryogenics relocated inside the chassis is a real achievement and is scored as one; it is not the elimination of a cooling requirement, and the closed-cycle cryocooler carries maintenance intervals, vibration, and failure modes that a wall-outlet system does not. Linear topology marked INFERRED.

Quantum Brilliance QB-QDK2.0 — undisclosed width

The most complete deployability story in the survey and the least complete performance disclosure. The system is real, operational, and independently hosted: installed at Fraunhofer IAF in May 2025 following a public tender, integrated directly into HPC infrastructure, with a prior unit at the Pawsey Supercomputing Centre. It fits a 19-inch rack, requires no cryogenics, and pairs QPU with NVIDIA GPUs and CPUs in one module. No qubit count, gate fidelity, or coherence figure for QDK2.0 appears in vendor press material or partner announcements. The only peer-reviewed figure available describes the prior-generation QB-QDK as a two-qubit room-temperature QPU. The ECW range is therefore left open at the upper end rather than estimated. This is the entry where right of reply is most likely to change the finding.

AQT IBEX Q1 and LYNX — throughput, not coherence

The strongest disclosure record in the survey and the entry where a common criticism must be rejected. Trapped ions are *not* coherence-limited: hyperfine coherence-to-gate-time ratios reach roughly 10^6, against roughly 10^3 for superconducting qubits. Any assessment attributing IBEX's limits to decoherence is wrong on the physics. All-to-all connectivity also means the routing tax scored against Equal1 is zero here, and that finding does not transfer.

The real constraint is absolute wall-clock. AQT publishes 12 fully connected qubits, up to 1000 gates per circuit, and more than 20,000 circuits per hour depending on depth. That last figure is the one buyers should read: a variational loop requiring millions of shots does not finish in useful time. Quantum Volume of 128 implies approximately 7 effective qubits, below the 12 physical, which is the honest width. LYNX publishes QV 32768, implying approximately 15 effective qubits — the widest figure in this survey and still less than half the lower floor band. LYNX's qubit count is not published; the launch material states explicitly that the advance is control and noise sensitivity rather than qubit count, which is an unusually candid framing and is recorded as such.

Specification discrepancy, logged. AQT's own product page reports average two-qubit error of (1.3 +/- 0.3)%, or 98.7% fidelity. Third-party listings report 97.7% and 98.85% for the same system. The vendor figure is used here; the discrepancy is recorded as an open item and is a candidate for the specification-drift log.

ORCA PT-1 and PT-2 — outside the width frame

Scored **NON-COMPARABLE** and deliberately not assigned an ECW. The PT series is a programmable boson sampler, not a universal gate machine: peer-reviewed description of the PT-1 architecture confirms a single photon source and single detector with programmable beam-splitters on a series of optical loops, shifting computation from the spatial to the time domain. Forcing a qubit-width number onto that architecture would produce a finding that is wrong rather than unfavourable, which is the failure mode this ledger exists to avoid. PT-1 is described at approximately 8 photonic qubits and PT-2 at

approximately 90 via increased temporal multiplexing; these are mode-and-photon counts and must not be tabulated against gate-model qubit counts. Deployment is genuine and multi-site, including PSNC Poland and the UK NQCC. The correct test for this system is problem-class match: whether the claimed workload lies inside the class the machine samples from, and whether a classical sampler reproduces the output distribution at the same accuracy. That test is not yet run and is listed as an open item.

Quandela Belenos — the multiplier is the state space

Twelve photonic qubits across 24 modes, gate-model, programmed through an open-source SDK, cloud-accessible to a large user base and delivered into EuroHPC infrastructure at CEA-TGCC. The deployment record is substantial and the engineering is real.

The headline performance multipliers require arithmetic. Belenos is published as approximately 4,000 times more computing power than the 6-qubit predecessor, and the 24-qubit Canopus as a 16-million-fold increase. Two raised to the twelfth is 4,096. Two raised to the twenty-fourth is 16,777,216. The multipliers are the dimension of the state space, not a measured performance result — and that same quantity is what sets the cost of simulating the machine classically. Recorded as a units finding under the addressability test, not as a disclosure failure: the qubit counts themselves are stated plainly and accurately throughout.

Two further items. Linear-optical gates are probabilistic, so effective width is a function of per-gate success probability and repetition count rather than register size; no success probability is published, which is why ECW is left undetermined rather than set to 12. And the cooling requirement is relocated, not removed: single-photon detectors and quantum-dot sources operate at cryogenic temperatures. That figure is graded COMMENTARY here, drawn from technical analysis rather than vendor documentation, and is an open item.

Five-category scorecard

Category	Grade	Finding
Hardware reality	**REAL — SCOPED**	All seven systems exist, ship, and rack-mount. Several are independently operated by national laboratories and supercomputing centres, which is stronger evidence than any vendor claim. What does not survive is width. Effective Computational Width across the entire surveyed field runs from 2 to approximately 15 qubits. The lower band of the classical simulation floor is 30. No system is within a factor of two of it.
Signal integrity	**PARTIAL — ASYMMETRIC**	Two of seven systems publish a two-qubit gate fidelity. AQT publishes the most complete record in the survey, including error bars and min/max across the register, and is scored accordingly. Equal1 publishes usable figures. The remainder publish partial figures or none. One vendor-versus-third-party discrepancy is logged. Where figures are absent, no value is inferred.
Post-processing dependence	**WHERE IT DIES**	Every system in the survey sits inside the range where exact classical simulation is instantaneous, which means no output requires the quantum hardware to have produced it. Architecture-specific costs land on the user: SWAP routing on constrained topologies, shot overhead on probabilistic photonic gates, repetition on sampling machines, and inter-core orchestration on multi-core designs. No vendor publishes a quantum-ablation baseline.
Threat relevance	**NONE — PROCUREMENT RISK**	No cryptographic implication anywhere in the survey and no movement on any PQC migration date. Breaking RSA-2048 requires thousands of logical qubits; the most favourable recent resource estimates put the physical requirement under a million at high fidelity. Two to fifteen physical qubits is not a point on that curve. The live risk is procurement: capital and headcount committed against a utility reading the hardware does not support.
Business decision	**SPLIT — REJECT CLAIM FUND TRAINING**	Reject the utility claim. The category is a training and integration purchase, and on those terms it is defensible: hybrid orchestration rehearsal, variational-loop engineering at real shots-per-second, operations experience with a stack that will matter later, and — for the sovereignty-constrained buyer — data residency that cloud access cannot provide. The honest alternative must be stated: at these widths a classical simulator teaches the same toolchain more cheaply and without hardware noise. What the hardware adds is the integration and operations rehearsal, not the qubits.

Kill conditions

The verdict is FALSE as stated. It is not permanent. Any of the following, demonstrated and documented, re-opens the scorecard and the affected entry is re-run rather than rewritten.

ECW above floor	A single addressable register above 30 qubits, with circuit structure stated. Between 30 and 50 the determination is a stated judgment against the specific circuit class, not a threshold result.
Composability	Where a headline is scored AGGREGATE: demonstrated entangling links between cores at fidelity sufficient to execute a circuit wider than one core, end to end, with the composed circuit and its fidelity published.
Fidelity at width	Two-qubit fidelity of 99.9% or better sustained across the full claimed register — not at best-case pairs — with coherence exceeding the depth the topology requires for an arbitrary pairwise gate.
Throughput	For trapped-ion entries: end-to-end wall-clock for a named workload at a stated accuracy target, including shot count, demonstrating completion in operationally useful time.
Problem-class match	For sampling architectures: a named workload inside the machine's sampling class where a current classical sampler fails to reproduce the output distribution at the same accuracy.
Quantum ablation	The condition no vendor in this survey currently satisfies, and the one that would settle the claim outright: a named production workload where the identical classical pipeline measurably degrades when the QPU is removed, with wall-clock decomposition and a best-classical comparison at the same accuracy target.
Disclosure	Publication of any figure recorded UNDISCLOSED, which re-opens the affected entry regardless of the value reported.

Open items and disclosure flags

Flag	System	Item
INFERRED	SAXON Q	Whether the sixteen cores have any inter-core coupling, or are fully independent. Taken from public commentary, not vendor documentation. The strength of the width finding depends on it.
INFERRED	Equal1	Strictly linear topology. The SWAP arithmetic assumes a linear chain; a different topology changes the routing tax and the lower end of the ECW range.
INFERRED	AQT	Effective qubit width derived from published Quantum Volume (QV 128 implies ~7; QV 32768 implies ~15). Derivation shown; not a vendor-stated width.
UNDISCLOSED	SAXON Q	Two-qubit gate fidelity. Recorded as a finding, not a gap. No value inferred.
UNDISCLOSED	Quantum Brilliance	Qubit count, gate fidelity, and coherence for QDK2.0. No performance figure for the current generation located in vendor or partner material.
UNDISCLOSED	AQT LYNX	Qubit count and gate fidelities. Quantum Volume published without the underlying register size.
UNDISCLOSED	Quandela	Per-gate success probability and heralding rate. Without these, effective width cannot be computed for a probabilistic linear-optical architecture and is left undetermined.
DISCREPANCY	AQT IBEX Q1	Vendor reports 98.7% average two-qubit fidelity; third-party listings report 97.7% and 98.85%. Vendor figure used. Candidate for the specification-drift log.
COMMENTARY	Quandela	Cryogenic operating temperatures for single-photon detectors and quantum-dot sources. Drawn from technical analysis rather than vendor documentation. Pending confirmation.
NOT RUN	ORCA	Problem-class match test. The correct assessment for a sampling architecture has not been performed and no verdict on utility is entered for this system beyond NON-COMPARABLE.
PENDING	All	Right of reply. Each vendor receives its rows and the findings concerning it before publication, with a stated response window. Corrections incorporated; responses printed verbatim; declining to respond recorded neutrally and without inference.

Status history

2026-08 · 005	Entry published. Supersedes the two-system evidence base of Scorecard 002 with a seven-system survey scored under the Hardware Claims Rubric v1.
2026-08 · 002	**Correction to a published kill condition.** Scorecard 002 set the classical-simulation threshold at 40 qubits addressable as one register. That threshold is clearable on paper by circuits that remain classically reproducible in fact: Clifford circuits are efficiently simulable at any width, and tensor-network methods run well past 40 at low entanglement. Replaced with the three-band criterion in this scorecard. The original grade stands; this is an addendum, not a replacement. The verdict on 002 is unaffected — both systems sat below either threshold.

The standing test

Score every rack-scale entrant on Effective Computational Width under its own binding constraint, published as a range with assumptions, and on whether the headline number names a quantity addressable in a single circuit. Not on the headline itself. A machine that advertises a total it cannot address as one register has reported a chassis inventory. A machine that advertises a multiplier equal to the dimension of its state space has reported how hard it is to simulate, which is not the same as what it computes.

PROVENANCE · SAXON Q SXQ128 and Equal1 RacQ public specifications, per Scorecard 002 · AQT product documentation for IBEX Q1 (12 qubits, QV 128, 1Q error (3.4 +/- 1.0)e-4, 2Q average error (1.3 +/- 0.3)%, crosstalk (1.0 +/- 0.4)%, up to 1000 gates, >20,000 circuits/hour) and LYNX Series launch material (QV 32768) · Amazon Braket provider documentation for IBEX Q1 · Quantum Brilliance and Fraunhofer IAF announcements for QB-QDK2.0; arXiv:2312.11673 for prior-generation QB-QDK architecture · ORCA Computing PT-1 and PT-2 product pages; arXiv:2409.13781 for PT-1 architecture · Quandela Belenos product page and launch material; Quandela Product Catalogue v6 · GlobalFoundries 22FDX · Gidney (2025) RSA-2048 resource estimates · Firebringer Quantum Hardware Claims Rubric v1 · All specification cells captured August 2026; archived copies held per rubric section 2.1.

METHOD · Effective Computational Width, per-modality instantiation, evidence grading, and the addressability test are defined in the Hardware Claims Rubric v1, published alongside this scorecard. Every figure in this document carries an evidence grade and a capture date. Where a finding rests on an INFERRED or COMMENTARY cell, it is listed as an open item and the verdict is scoped accordingly.