Rack-scale quantum utility

A claim-level assessment of the deployability-versus-utility conflation in commercially available rack-mounted quantum systems. Graded against the five-category scorecard. Published August 2026.

> **“Commercially available rack-scale quantum systems are useful for production workloads today.”**

CLAIM CLASS	Market-interpretation claim. Graded as the generic formulation circulating in buyer conversations, not as a review of any named vendor or product.
VERDICT	**FALSE — as stated.** Deployability is real and achieved. Utility is not. The two are being read as one claim, and the purchase decision is being made on the wrong one.
EVIDENCE BASE	SAXON Q SXQ128 (NV-centre diamond, room temperature, wall outlet, 16 × 8-qubit cores) and Equal1 RacQ (silicon spin quantum dots, 0.3 K on-chip cryogenics, 6 qubits, GlobalFoundries 22FDX).
REUSABLE TEST	Maximum coherent circuit width and inter-core coupling — not headline qubit count. Any future rack-scale entrant is scored on the same two variables.

What is actually being claimed

Two rack-mounted quantum systems reached general availability with different thermal architectures and near-identical facilities profiles. The market has read that as a utility milestone. It is a deployment milestone.

The common correction — disputing “room temperature” on the grounds that on-chip cryogenics is still cryogenics — is a category error, not a finding. From a facilities standpoint both machines solve the same problem: no cold room, no custom infrastructure, no dedicated lab. Thermal architecture is not the dividing line. The dividing line is between **deployability** and **utility**, and both vendors have delivered the first while the market is hearing the second.

One precision that must not be flattened: eliminating the facility requirement is not the same as imposing an equal operational burden. RacQ is 400 kg and 1.6 kW, requires reinforced flooring and a dedicated 20 A circuit, and contains a closed-cycle cryocooler — a mechanical system with maintenance intervals, vibration, and failure modes. SXQ128 is a wall outlet. Both eliminate the facility requirement; they do not impose the same operational burden.

Five-category scorecard

CATEGORY	GRADE	FINDING
Hardware reality	REAL — SCOPED	The machines exist, ship, and rack-mount. That part of the claim survives. What does not survive is the width: the largest circuit a user can address as one coherent register is **8 qubits** on SXQ128 and **6** on RacQ. The advertised 128 is sixteen disconnected eights. Facilities requirement: genuinely eliminated by both. Operational burden: not equivalent (see above).

CATEGORY	GRADE	FINDING
Signal integrity	PARTIAL — ASYMMETRIC	Equal1 publishes two-qubit gate fidelity: 99.3% at 200 ns, interleaved. SAXON publishes 99.92% single-qubit and omits two-qubit entirely. These cannot be scored on identical criteria and no number is inferred for the omission — but **undisclosed two-qubit fidelity is itself informative**, and is recorded as a finding rather than a gap. On the published figure: 99.3% clears the ~99% surface-code threshold, but clearing threshold is not the same as affording it. Barely above threshold means punishing physical-to-logical overhead.
Post-processing dependence	WHERE IT DIES	Neither headline number is addressable as a single register, and the routing cost is carried by the user. On a linear 6-qubit chain, entangling the end qubits costs four SWAPs — twelve CNOTs of tax before the intended gate fires. $0.993^{13} \approx 91.2\%$, and a 200 ns gate becomes roughly 2.6 µs of accumulated noise. On the multi-core side, parallel sub-circuits are not a workaround; they are a different machine model. Independent registers give **throughput, not width**: N separate 8-qubit circuits is not one 8N-qubit circuit, and no scheduler closes that gap without entanglement between cores. Add inter-core transfer and you are describing a distributed quantum architecture — an open research problem with its own overhead, not a feature you get free by putting chips in one chassis.
Threat relevance	NONE — PROCUREMENT RISK	No cryptographic implication. Breaking RSA-2048 requires thousands of logical qubits; even the most favourable recent resource estimates put the physical requirement under a million at high fidelity. Six to eight physical qubits is not a point on that curve, and neither machine moves a PQC migration date. The live risk is procurement, not cryptography: capital and headcount committed against a utility claim the hardware does not support.
Business decision	SPLIT — REJECT CLAIM FUND TRAINING	**Reject "useful for production workloads." Fund the toolchain.** This is a training purchase, not a compute purchase — and as a training purchase it is defensible. What you are buying is hybrid orchestration rehearsal, variational-loop engineering at real shots-per-second, and staff fluency in a stack that will matter later. What you are not buying is a workload. Note the honest alternative: at 6 to 8 qubits a classical simulator teaches the same toolchain more cheaply, minus the hardware noise. The residual justification for the hardware is the integration and operations rehearsal — the quantum-placeholder step in a classical-hybrid workflow — not the qubits.

The utility floor, stated as a criterion

"Eight qubits is not much use, six less so" is a judgement. Stated as a criterion it is a test, and the conclusion follows from the test rather than from the author:

> **Criterion:** any circuit width that can be exactly simulated on commodity hardware cannot deliver value the classical machine cannot already deliver. The utility floor therefore sits wherever exact classical statevector simulation stops being trivial.

At six to eight qubits you are at $2^6 = 64$ to $2^8 = 256$ amplitudes. That is not merely simulable; it is instantaneous on a laptop, and on a phone. The floor is not crossed at these widths by any argument about algorithms, because there is no algorithm whose output the classical machine cannot reproduce exactly and faster.

This also disposes of the pedagogical defence. If the machine is for learning, a classical simulator teaches the same toolchain more cheaply and without the hardware noise. What the hardware adds is the integration rehearsal — scheduling, orchestration, the physical peer-accelerator pattern — which is a real thing to buy, and a different thing from

the claim under test.

Qubit count is not the ultimate tell of utility. But a floor exists below which nothing is possible, and both machines are under it.

Evidence base, scored on the honest metric

	SAXON Q SXQ128	Equal1 RacQ
Headline number	128 qubits	6 qubits
Max coherent circuit width	8	6
Topology	16 independent 8-qubit cores	Linear chain, SWAP routing
Inter-core coupling	None (INFERRED — unconfirmed)	N/A — single register
1Q gate fidelity	99.92%	99.9% @ 140 ns
2Q gate fidelity	UNDISCLOSED	99.3% @ 200 ns, interleaved
Readout fidelity	Not published	99% @ 10 µs, parity
Thermal architecture	Room temperature, NV diamond	0.3 K on-chip cryo, silicon spin
Facility requirement	None — wall outlet	None — self-contained rack
Operational burden	Low	400 kg, 1.6 kW, 20 A, cryocooler service

Scoring note: the two machines are not scored on identical criteria, because the disclosures are not identical. Where a figure is undisclosed it is marked undisclosed and no value is inferred. Where a structural property is taken from public commentary rather than vendor documentation it is marked INFERRED and carries an open item below.

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 entry is re-run rather than rewritten:

— **Inter-core entangling links.** Demonstrated coupling between cores at fidelity sufficient to compose registers — i.e. a circuit of width greater than one core, executed end-to-end. This is the single largest lever on the multi-core architecture and would move the width finding immediately.

— **Maximum coherent circuit width above the classical-simulation floor.** A single addressable register wide enough that exact statevector simulation is no longer trivial. The threshold used here: **40 qubits addressable as one register**, above which exact simulation leaves commodity hardware.

— **Two-qubit fidelity of 99.9% or better, sustained at that width**, with coherence time exceeding the SWAP depth the topology requires for an arbitrary pairwise gate. Fidelity at nearest-neighbour pairs alone does not satisfy this.

— **A quantum-ablation baseline on a production workload.** A named workload where the identical classical pipeline demonstrably degrades when the QPU is removed, with wall-clock decomposition and a best-classical comparison at the same accuracy target.

— **Disclosure of the withheld figures.** Publication of SAXON two-qubit gate fidelity changes the asymmetry finding regardless of the value reported.

Open items and disclosure flags

INFERRED	Whether SAXON's sixteen cores have any inter-core coupling, or are fully independent. Taken from public commentary, not vendor documentation. Pending confirmation; the strength of the width finding depends on it.
INFERRED	Whether Equal1's six qubits are strictly linear. The SWAP arithmetic above assumes a linear chain; a different topology changes the routing tax.
UNDISCLOSED	SAXON two-qubit gate fidelity. Recorded as a finding, not a gap. No value inferred.
SCOPE	This scorecard grades the market claim. It is not a product review and does not assert intent on the part of either vendor. Both machines are accurately described by their own published specifications; the failure is in the interpretation the market has placed on them.

The standing test

> Score every rack-scale entrant on **maximum coherent circuit width** and **inter-core coupling**. Not on the headline number. A machine that advertises a total and cannot address it as one register has reported a chassis inventory, not a processor.

PROVENANCE · SAXON Q SXQ128 and Equal1 RacQ public specifications and technology pages · Equal1 published qubit performance figures (1Q 99.9% @ 140 ns randomized Clifford; 2Q 99.3% @ 200 ns interleaved; readout 99% @ 10 µs parity) · GlobalFoundries 22FDX process · Dell PowerEdge R770 integration demonstration · Gidney (2025) RSA-2048 resource estimates · Firebringer Quantum teardowns, August 2026