FOR IMMEDIATE RELEASE

RunIQ Labs · press@runiqlabs.ai · www.runiqlabs.ai

RunIQ Labs Launches to Fix the AI Accountability Gap

The Measurement Crisis Costing Enterprises Millions in Ghost Capacity™

RunIQ Labs Launches as the Independent Measurement Authority for AI Infrastructure

SAN DIEGO, CA — May 12, 2026 — Enterprise AI infrastructure has scaled faster than its ability to be measured. RunIQ Labs today emerged from stealth to address the AI Accountability Gap — the widening divide between what enterprises pay for AI infrastructure and what that infrastructure actually delivers. Concurrent with its launch, the company released a landmark whitepaper, "What Really Happens When You Run AI at Scale: How Ghost Capacity™ Impacts the Bottom Line," documenting that peak benchmarks overstate real-world production performance by up to 39% — and that the wasted spend hiding in that gap now has a name: Ghost Capacity™.

Unlike vendor benchmarks and synthetic lab tests, RunIQ Labs executes real AI workloads on real infrastructure — producing measurement that is immediately actionable at the deployment level. The platform shows enterprises what their AI is actually delivering in production and where they are wasting money. That is the question every CIO and CFO needs answered. RunIQ was built to answer it.

Ghost Capacity™ is provisioned, paid-for AI infrastructure that is not delivering economically useful work. It is the direct consequence of a market where every vendor has a tool — and none of them have the structural incentive to tell enterprises the full picture. RunIQ Labs' research found that configuration alone can account for a 22% throughput variance on identical hardware — and that short-duration benchmarks inflate efficiency metrics by up to 39% versus sustained production reality. It is not a hardware failure. It is a measurement failure. RunIQ was built to fix it.

> *"Every vendor in the AI stack has a tool. None of them have the structural incentive to tell you the full picture. Model correctness does not guarantee production efficiency. Our data shows that configuration alone can cost an enterprise 22 percent of its throughput — with no changes to hardware, model, or precision. That is not a technology failure. It is a measurement failure. RunIQScore™ was built to fix it, independent of vendor influence."*
>
> **— RunIQ Labs**

The PPC Standard: Performance, Power, and Cost

RunIQScore™ moves beyond raw throughput to unify Performance, Power, and Cost into a single intelligence layer. Derived from sustained, production-style workloads — not peak-only lab figures — the platform provides:

- True Economic ROI: Measuring Joules-per-token and actual cost-per-inference to identify the most efficient hardware for specific model architectures
- Architecture and Hybrid Validation: Benchmarking sustained performance across high-performance accelerators and cloud-native silicon to ensure cross-platform parity
- Optimisation Guardrails: Identifying saw-tooth under-utilisation patterns that signal infrastructure waste in both public cloud and private data centres

Together these measurements produce RunIQ's Total Cost of Intelligence (TCIO) — a deployment-grade economic model that makes AI spend visible, comparable, and defensible. Every AI decision is an economic decision. TCIO makes it measurable. Unlike benchmark standards requiring industry-wide consensus, RunIQScore™ generates immediate deployment-level value from day one — with credibility compounding as DDR evidence accumulates across enterprise deployments over time.

The platform generates Deployment Diagnostic Records (DDR) — cryptographically sealed, audit-ready documentation of infrastructure configuration, workload behaviour, and measurement methodology. For boards and regulators demanding defensible evidence of AI system behaviour — not just dashboards — the DDR provides the audit-grade record the market has lacked. For organisations operating under EU AI Act, NSPM-8, or sovereign AI mandates, DDRs provide independently generated evidence of AI infrastructure behaviour produced by a platform with no commercial interest in the result.

The Creeper Syndrome: Why Continuous Measurement Is Not Optional

"Drift is inevitable. Detection is optional."

RunIQ Labs also introduced RunIQMonitor, the continuous measurement layer that detects Creeper Syndrome — the natural, gradual erosion of AI deployment efficiency over time. Without continuous measurement against a verified baseline, enterprises risk eroding their effective performance baseline silently across billing cycles. RunIQMonitor turns invisible drift into controlled outcomes.

FIND. RECOVER. DEFEND.

RunIQ's three-step engagement model provides a complete lifecycle for AI infrastructure integrity. The DDR Assessment (FIND) provides independent, cryptographically verifiable measurement of Ghost Capacity™. The recovery phase converts identified inefficiency directly into infrastructure return — no new hardware required. RunIQMonitor (DEFEND) then sustains that return against the inevitable drift that follows every model update, workload shift, and infrastructure change.

Strategic Ecosystem Validation

The launch follows a period of intense technical validation with the pillars of the AI supply chain. The AI processor market has expanded from two viable GPU suppliers in 2014 to 161 companies today — spanning commercially shipping products and active development across seven categories. This fragmentation creates measurement gaps no single vendor is positioned to close, and establishes why an independent measurement authority is a structural necessity, not a preference.

RunIQ Labs's platform is currently undergoing deep-stack software assessment with the engineering leadership of a Tier-1 GPU vendor and is finalising a three-way Proof of Concept with a leading Global Systems Integrator and a Tier-1 Hyperscale Cloud Provider to validate enterprise-scale agentic workloads.

> *"The AI processor market has grown from two viable GPU suppliers in 2014 to 161 companies today — spanning commercially shipping products and active development across seven categories. That fragmentation creates measurement gaps no single vendor is positioned to close. RunIQ is doing the work the market structure requires."*
>
> **— Jon Peddie, President, Jon Peddie Research**

> *"The AI infrastructure market has arrived at a structural accountability gap. Enterprise AI spend has grown rapidly, but the tools available to measure return on that spend remain fragmented, vendor-biased, or inadequate for production conditions. RunIQ Labs is one of the first companies attempting to address this gap systematically, with a platform aligned to the increasing demand for measurable outcomes, energy-aware infrastructure decisions, and defensible evidence of AI system behavior in production environments."*
>
> **— Dave McCarthy, Research Vice President, IDC**

RunIQ Partner Programme

RunIQ Labs also announced its Partner Programme, offering Foundation, Strategic, and Premier tiers. MSPs and System Integrators in the programme receive RunIQ's Verification Receipt — the independent proof layer that enables partners to demonstrate, protect, and monetise AI deployment value for their clients. RunIQ is the Intel Inside for AI trust.

About RunIQ Labs

RunIQ Labs is the independent measurement authority for enterprise AI infrastructure — the AI Integrity Layer that delivers the one answer enterprises actually need: is your AI working, is it efficient, and is it worth what you are paying for it?

Through the RunIQScore™ framework, RunIQ Labs unifies performance, power, and cost into a deployment-grade intelligence fabric. The platform generates Deployment Diagnostic Records and Total Cost of Intelligence analysis — providing audit-ready, independently generated evidence of AI infrastructure performance across the full deployment lifecycle.

Qualified organisations can request a complimentary DDR Assessment at contact@runiqlabs.ai. Selected organisations will receive a Deployment Diagnostic Record from their production infrastructure within 72 hours. For more information and to download the whitepaper, visit www.runiqlabs.ai.

Media Contact

RunIQ Labs

press@runiqlabs.ai · www.runiqlabs.ai

Note: *Ghost Capacity™ and RunIQScore™ are trademarks of RunIQ Labs. The 22% throughput variance and 39% benchmark inflation figures are derived from sustained production workloads across multiple real-world deployments. Detailed methodology is available in the company's whitepaper.*

###