

Independent Third-Party Verification Proves NOL8 Delivers Up to 118x More Data Per Dollar Than Google RE2

Testing now proves NOL8's economic advantage over Google RE2 at enterprise-scale complexity and data volumes.

- **Independent third-party verification** — Professor Amit Klein, Head of the Federmann Cyber Security Center at The Hebrew University of Jerusalem, independently verified NOL8's performance against Google RE2.
- **Up to 118x more data per dollar** — NOL8 delivered between 4.1x and 118x more data per dollar of infrastructure cost than Google RE2, with its advantage widening at every step up in workload complexity.
- **Testing Proves Enterprises no longer need to trade governance for lower cost** – As scrutiny of AI governance grows, the testing proves enterprises no longer have to trade tighter data protection for higher compute costs.
- **Cost held flat as complexity rose** — NOL8's cost to process data held at approximately US\$2 per GB/s per hour across every test, while RE2's rose from US\$8 to US\$240.
- **Latency held at 3–5 milliseconds** — NOL8 held latency at 3–5 milliseconds up to 98% of the workload while delivering over 950 MB/s of throughput. Google RE2 latency rose to 10 milliseconds and its throughput fell to 2 MB/s.
- **Accuracy, scalability and Economics proven for channel partners** — the results strengthen NOL8's commercial case across the Megaport (ASX: MP1) and Equinix (NASDAQ: EQIX) networks, where customers pay for every byte of data processed.
- **World's first fully deterministic data processing platform** — NOL8 inspects, filters and applies security and governance rules to every byte of data in motion, with identical outcomes every time, regardless of complexity or load allowing enterprise grade scalability.
- **Class C milestone achieved** — the independent verification has satisfied the Vesting Condition for the Company's Class C Performance Rights.

FortifAI Limited (ASX: FTI) ("FTI" or the "Company") is pleased to announce the results of independent, third-party verification of its NOL8 AI Data Plane against Google RE2, a widely adopted engine for data pattern matching at speed and scale. In real-world application testing, NOL8 delivered between 4.1x and 118x more data per dollar than RE2, with the advantage widening as workload complexity increased.

Enterprise AI is about putting data in and getting intelligence out. With growing scrutiny over enterprise data sensitivity, every byte should be inspected, filtered and governed before it enters an AI workload, and the result must be correct every time. NOL8 is designed as the world's first fully deterministic data processing platform, delivering byte-for-byte accuracy at the lowest latency, highest throughput and lowest cost. This successful testing measured the economics and scalability of that capability.

Independent Verification

NOL8 engaged Professor Amit Klein, Head of the Federmann Cyber Security Center at The Hebrew University of Jerusalem, to independently verify the capability of NOL8 Technology V1.0 to deliver improved data per dollar against industry-standard CPU pattern matching. Professor Klein's assessment was based on his independent review of NOL8's testing methodology and results. The testing involved pushing an increasing number of rules (128 to 5,000 rules) at chaotic intervals. The use of chaotic load distribution reflects peaks and idle time in real world load scenarios. The common denominator for testing was the use of AWS with NOL8 Technology V1.0 utilising FPGA compared to Google RE2 utilising CPUs. Google RE2 is a widely adopted engine for data pattern matching at speed and scale.

NOL8 was tested against Google RE2 running on CPUs. Six tests were run on both systems. Each test increased the number of rules, measured latency and throughput, and divided the result by hardware cost to derive a data-per-dollar figure.

Rules set how tight the filter's mesh is. The tighter the mesh, the more it catches. More rules give enterprises greater certainty that sensitive data is not passing into their AI systems. It is also where software-based filtering begins to break down.

Latency and Throughput Advantages Further Reinforced

NOL8 held latency between 3 and 5 milliseconds up to 98% of the workload, while delivering over 950 megabytes per second of throughput.

RE2 latency rose quickly to 10 milliseconds. Over the same tests, RE2 throughput fell from approximately 100 megabytes per second to only 2 megabytes per second, against NOL8's 950 megabytes per second.

Full, deterministic inspection of every byte no longer comes at growing infrastructure cost.

At the smallest ruleset of 128 rules, NOL8 delivered 4.1x more data per dollar than RE2. NOL8's advantage grew as rule complexity increased. At the largest ruleset of 5,000 rules, NOL8 delivered 118x more data per dollar.

NOL8's cost held flat at approximately US\$2 per GB/s per hour across all six tests. RE2's cost rose thirty-fold, from US\$8 to US\$240. The breakthrough test proves that enterprises must no longer trade tighter data protection against rising infrastructure cost.

Google RE2 costs climb as complexity grows. Nol8 costs remain flat.



Source: Independent verification by Professor Amit Klein, Federmann Cyber Security Center, The Hebrew University of Jerusalem. Cost shown in US\$ per GB/s of throughput per hour.

What It Means for NOL8's Megaport and Equinix Distribution Network

NOL8's neo-cloud distribution model allows the technology to reach enterprise AI buyers through the networks enterprises already use.

In June 2026, NOL8 joined the Megaport AI Exchange alongside Megaport's compute, network and storage layers for proof-of-concept deployments.

In August 2026, multi-continent testing across Megaport's private network in Virginia, Tokyo and Sydney delivered approximately 3 milliseconds of in-region latency and 10x more governed data per dollar than legacy CPU-based software. For clarity, the purpose of the Megaport test was to prove 100% byte coverage across a global network, not to stress-test the system with increasing rule count as this independently-verified benchmark did. Taken together, the tests supported NOL8's claim of its ability to filter 100% of enterprise data headed to AI workloads, with cost staying flat as load increases, while the data is sent back and forth across a global network to any neocloud in the world.

The partner network was bolstered in September 2026, where NOL8 became available through Equinix Fabric across more than 60 locations globally.

Accuracy is the foundation of trust in AI governance. A platform that returns the same answer, byte for byte, on any network and at any scale gives partners a product they can stand behind.

Both the Megaport and Equinix channels sell the same outcome: governed data, delivered to AI workloads at the lowest cost. For NeoClouds and GPU providers on the Megaport AI Exchange, a data processing technology whose cost stays flat as complexity grows means more filtered, governed data reaching every GPU for each dollar spent. For enterprises connecting clouds, AI providers and their own infrastructure over Megaport and Equinix networks, governance policies can tighten as AI use expands without infrastructure cost compounding.

Independent verification gives Megaport, Equinix and their customers a third-party basis on which to assess NOL8 through current proof-of-concept engagements ahead of general availability in the coming weeks.

Milestones Delivered Ahead of Schedule

The results extend NOL8's published performance benchmarks with Google RE2 announced in April 2026. Further, the NOL8 February 2026 product roadmap set the path to commercialisation. These results provide further third-party verification ahead of general product availability by late-October, ahead of NOL8's forecast timeline.

"Every enterprise deploying AI faces the same trade-off. The more tightly it protects its data, the more it pays in compute. Independent testing shows NOL8 removes that trade-off where our cost holds flat as protection tightens, while CPU-based software system costs climb thirty-fold, leaving us with a 118x data-per-dollar advantage at the highest level of complexity tested. NOL8 inspects every byte and applies every rule with the same result every time. That determinism is what enterprises need to trust AI with their data.

Third-party verification of that result matters to our partners at Megaport and Equinix, and to the customers they serve. It gives them a proven economic case to take NOL8 to market."

— Kelly Herrell, Chief Executive Officer, FortifAI

Vesting of Performance Rights

The independent verification of NOL8's capability on a data-per-dollar basis against an industry-standard CPU pattern matching engine satisfies the vesting condition for 52,516,667 Class C Performance Rights. The Class C Performance Rights were issued pursuant to the Company's acquisition of FastAI Pty Ltd first announced on 15 December 2025, and shareholder-approved CEO equity incentives.

12 October 2026

Authorised for release by the Board of Directors of FortifAI Limited.

Media Contact

Matt Wolpin | matt@nol8.io

Investor Relations

John Marrin | john.marrin.ir@nol8.io | +61 468 828 878

Forward-Looking Statements: This announcement may contain forward-looking statements that involve risks and uncertainties. Actual results may differ materially from those anticipated. FortifAI undertakes no obligation to update forward-looking statements.

About NOL8

AI systems make decisions on data that teams often can't fully see or verify in the moment, and NOL8 exists to close that gap. Its technology, the AI Data Plane, gives organisations visibility into how data moves through their AI systems, enabling certainty where there was once assumption. For more information, visit nol8.io.

About FortifAI

FortifAI Limited (ASX: FTI) is an Australian Securities Exchange listed technology company focused on artificial intelligence and advanced computing through its subsidiary NOL8.