
Benchmark Testing Against Google RE2 Delivers World-First Results

Confirming more than a 200,000x Throughput Advantage

Highlights

- **Nol8's world-first technology delivers a throughput advantage of over 200,000x under AI-grade workloads against Google's RE2 Engine** — a widely adopted engine for data pattern matching at speed and scale.
- **Nol8's FPGA-accelerated engine maintains a constant 1,500 MB/s throughput** across all complexity tiers and all load conditions — while the RE2 software collapses to near-zero as complexity and load increases.
- **The results confirm that Nol8 has unlocked a previous ceiling to data processing scalability** — a barrier no software solution has been able to overcome.
- The testing provides further validation as the Company closes-in on its **Enterprise Ready Benchmarking Engine in June 2026** (announced to the ASX on 16 February 2026).
- Results published represent only a portion of Nol8's total performance capability — **future benchmark testing will be conducted** and published as the Company continues to explore the full limits of its world-first technology.
- Benchmark **data throughput results** have been published to date — **Work is ongoing to quantify these gains as economic reductions in hardware footprint**, computational load and infrastructure dependency.

FortifAI Limited (ASX: FTI) ("**FTI**" or the "**Company**") is pleased to announce the results of benchmarking of its Nol8 AI Data Plane technology against Google's RE2 engine — the globally accepted standard for data pattern-matching at high-performance.

The results confirm that **Nol8's neural-network-based algorithm, hyper-accelerated by FPGA hardware, has unlocked a previous ceiling to data processing scalability** that no software-based architecture has been able to surpass.

About Nol8's Technology

Nol8 is using its technology to establish a new category in AI infrastructure: the AI Data Plane — the high-speed infrastructure layer that sits between raw data and AI inference, purpose-built to solve the data processing constraints of unstructured AI workloads. Built on algorithmically enhanced Longest Prefix Matching, scaled by machine learning neural networks, and hyper-accelerated by FPGA hardware, Nol8's engine processes data-in-motion at millisecond-grade latency without buffering or batching.

Backed by five years of published academic research by its founders, Nol8 has unlocked a previous ceiling to data processing scalability that no software-based architecture has been able to surpass. Nol8 conducted this exercise ahead of the planned release of its enterprise ready benchmarking engine by June 2026.

Unlike CPU-based software approaches, Nol8's FPGA architecture processes data in parallel at the hardware level. Performance does not degrade as workload complexity increases. It is constant, predictable, and unlike any solution previously available.

Why FTI Conducted This Benchmark Test

As the global AI industry accelerates toward agentic and autonomous systems, the volume and complexity of data that must be processed in real time is growing exponentially. Enterprises deploying AI at scale — across cybersecurity, financial services, healthcare, and AI infrastructure — require data processing systems capable of handling thousands of simultaneous classification rules without performance degradation.

Nol8 conducted benchmarking against Google RE2 to provide an objective, verifiable measure of its technology's performance against the most widely adopted software standard. The test was designed to reflect real-world enterprise AI workloads across three complexity tiers — from simple web filtering through to the full demands of agentic AI data pipelines.

"The purpose of this benchmark was simple: to put our technology against the best available software standard and let the results speak for themselves. What these numbers confirm is that we have crossed a threshold that software alone cannot surpass.

*The purpose of the exercise was to establish a credible reference case for how software-based pattern matching performs as rule complexity and load increase, and to assess Nol8's expected performance advantage under those conditions. **The scalability ceiling that has constrained AI data infrastructure is not a software problem — it is an architectural one. Nol8 solves it at the hardware level.**"*

— Alon Rashelbach, Co-Founder and CTO, Nol8

Benchmark Results: Nol8 vs Google RE2 — Throughput (MB/s)*

All results measured in megabytes per second (MB/s). Nol8's result is constant at 1,500 MB/s across all complexity tiers and all load conditions.

| COMPLEXITY TIER | P50 (RE2) | P90 (RE2) | P99 (RE2) | NOL8 (ALL TIERS) | NOL8 P99 ADVANTAGE |
|---|-----------|-----------|------------|------------------|--------------------|
| Low (~10 rules) Basic Web / API Applications | 46 MB/s | 22 MB/s | 10.8 MB/s | 1,500 MB/s | ~138x |
| Medium (~1,000 rules) Enterprise security | 4.5 MB/s | 1.7 MB/s | 0.28 MB/s | 1,500 MB/s | ~5,400x |
| High (6,000+ rules) AI-grade data classification | 1.5 MB/s | 0.43 MB/s | 0.007 MB/s | 1,500 MB/s | ~200,000x |

*See below explanation on Load Percentiles and Complexity Tier

Source: Company commissioned RE2 benchmarking. RE2 results derived from standard benchmarking methodology. Nol8 performance measured on FPGA hardware appliance.

The AI Data Plane Opportunity

The global datasphere is forecast to grow from 334 Zeta Bytes in 2025 to 19,267 Zeta Bytes by 2035, driven overwhelmingly by unstructured data generated by AI agents, autonomous systems, and large language models operating at continuous scale. 90% of future data flows will be unstructured — the most compute-intensive category of data to process. (refer to company ASX presentation: *AI Data Plane* – 17th February 2026)

This data must be filtered, enriched, classified, and routed in real time — before it reaches the model. This is the AI Data Plane: the missing infrastructure layer that the industry now requires. Nol8's benchmark results confirm it has built this layer, and that its performance is in a category of its own.

"What we are seeing in these results is the practical consequence of a genuine architectural breakthrough. Software has been reaching its limits for years — enterprises have been papering over this with CPU arrays that are costly, complex, and energy-hungry. This benchmark proves that there is a better way. The AI era demands infrastructure that was built for it. That is what Nol8 has created."

— Alon Rashelbach, Co-Founder and CTO, Nol8

Further Optimisation and Infrastructure Load Testing Results To come

The results published in this press release represent only a portion of the technology's total performance capability. Further benchmark testing will be conducted and results published as the Company continues to explore the full limits of its world-class technology.

To date, the Company has released throughput performance results only. Work is ongoing to further translate these gains into measurable infrastructure reduction outcomes — specifically, reducing computational load and hardware dependency while maintaining stable low latency at scale.

Understanding P50, P90 and P99: What the Load Percentiles Mean

The benchmark measures performance at three load scenarios — P50, P90, and P99. An example scenario in the context of Amazon has been used.

| LOAD CONDITION | WHAT IT MEANS | AMAZON ANALOGY |
|--------------------|---|---|
| P50 — Median | Normal operating conditions — the experience half the time | A regular Tuesday on Amazon. Normal traffic, predictable load, no surprises. |
| P90 — High Load | Top 10% of traffic events — elevated but not exceptional pressure | The lead-up to Christmas. Traffic is surging. The system is under sustained, elevated load. |
| P99 — Extreme Load | The worst 1% of all events — maximum real-world stress | Black Friday. Every shopper simultaneously using the service. The entire system under its highest conceivable pressure at once. |

P99 is the figure that matters most for AI infrastructure. Agentic AI systems operate continuously — 24 hours a day, 7 days a week. Traffic spikes are not anomalies to be planned around. They are operating conditions to be built for. A system that collapses at P99 is not enterprise-grade. At P99 and 6,000+ rules — the conditions that define real-world AI data classification — RE2 chokes, allowing just 0.007 MB/s of traffic. NoI8 allows 1,500 MB/s of traffic throughput. The gap is over 200,000x.

Understanding The Complexity Tiers

Rather than technical rule counts, think of complexity as the sophistication of the task being performed on data. Three tiers were tested:

| TIER | COMPLEXITY | TYPICAL REAL-WORLD SCENARIO | RULE COUNT |
|--------|------------|--|--------------|
| Low | Basic | Static websites, simple API gateway filtering, standard AI security — the scenarios AWS Bedrock Guardrails were designed for | ~10 rules |
| Medium | Enterprise | Corporate threat detection, SIEM platforms, log analysis, fraud screening — the backbone of enterprise security infrastructure | ~1,000 rules |
| High | AI-Grade | Real-time agentic AI pipelines, LLM data classification, unstructured data governance, multi-jurisdiction compliance at scale | 6,000+ rules |

At Low complexity, RE2-based software performs adequately — it was designed for this. At High complexity, enterprises today compensate by deploying large arrays of CPUs working in parallel: thousands of processors to spread the load and reduce latency. This approach is expensive, power-intensive, and structurally unable to keep pace with exponentially growing AI data volumes. NoI8 eliminates this requirement entirely.

What is Google RE2 — And Why Is It the Benchmark?

RE2 is Google's open-source regular expression (regex) engine, widely regarded as one of the fastest and safest software-based data pattern-matching engines available. It is used across Google's own infrastructure and in widely deployed modern infrastructure software and enterprise data platforms. Its primary advantage over alternative software engines is a guaranteed linear processing time — meaning performance degrades predictably rather than catastrophically under adversarial or complex inputs.

Separately, Amazon's Bedrock Guardrails — AWS's safeguard layer for generative AI applications — publicly documents a limit of up to 10 custom regex rules in its sensitive-information policy configuration. This illustrates how regex-based rule systems can run into practical limits as AI workloads grow in scale and complexity.

RE2 is a widely used software reference point for regex performance and safety precisely because it represents the ceiling of what software can achieve. NoI8's benchmarking against RE2 is therefore a direct measure of where hardware-accelerated neural network processing stands relative to the best software alternative available today.

Authorised for release by the board of FortifAI Limited.

Media and Public Relations contact:

Emily Walkerden (New York)

emily@nol8.io

+1 (917) 545 5817

About The Nol8 Technology

Nol8 is building the foundational AI Data Plane for the era of Autonomous Agents. By combining Neural-Network-Based Algorithms with FPGA hardware acceleration, Nol8 delivers unprecedented speed, efficiency, and scale for the world's most demanding AI data environments. Nol8's world-first technology has unlocked a previous ceiling to data processing and scalability — enabling enterprise AI systems to operate at the speed of the data stream itself.

About FortifAI

FortifAI Ltd (ASX: FTI) is an ASX listed AI infrastructure company developing and commercialising technology with a focus on AI. In addition to developing AI-forward technologies, FortifAI has developed a broad portfolio of video games. FortifAI uses AI to target efficiencies and expansion opportunities in technology.

Glossary

The following company and industry terms are used in this announcement:

- **Agentic AI:** AI systems that operate autonomously and continuously, executing decisions and actions in real time without human input at each step.
- **AWS Bedrock Guardrails:** Amazon's managed safeguard layer for generative AI applications, which uses RE2-based pattern matching to filter and govern data flows into and out of foundation models.
- **CPUs:** Central Processing Units — the primary general-purpose processors used in conventional computing infrastructure.
- **FPGA:** Field-Programmable Gate Array — a specialised semiconductor device whose hardware logic can be configured after manufacture, enabling parallel, high-speed data processing that far exceeds the throughput of conventional CPUs.
- **LLM:** Large Language Model — an AI model trained on large volumes of data to understand and generate human language.
- **MB/s:** Megabytes per second — a measure of data throughput indicating how much data a system can process in one second.
- **P50:** The median performance result — representing system behaviour under normal, typical operating conditions.
- **P90:** The 90th percentile performance result — representing system behaviour under elevated, high-load operating conditions experienced in the top 10% of traffic events.
- **P99:** The 99th percentile performance result — representing system behaviour under extreme load, reflecting the most demanding 1% of real-world operating conditions.
- **RE2 or RE2-based software:** Google's open-source regular expression (regex) engine — the globally accepted benchmark for data pattern-matching performance.
- **SIEM platform:** Security Information and Event Management platform — enterprise software that aggregates and analyses security event data in real time for threat detection and compliance.
- **ZB:** Zettabyte — a unit of digital data equal to one trillion gigabytes, used to measure global data volumes at scale.

Forward Looking Statement and Disclaimer

Certain statements contained in this ASX release, including information as to the future financial or operating performance of the Company and its projects, may be forward looking statements. Such forward looking statements:

- (a) are necessarily based upon a number of estimates and assumptions that, while considered reasonable by the Company, are inherently subject to significant technical, business, economic, competitive, political and social uncertainties and contingencies;*
- (b) involve known and unknown risks and uncertainties that could cause actual events or results to differ materially from estimated or anticipated events or results reflected in such forward looking statements; and*
- (c) may include, among other things, statements regarding estimates and assumptions in respect of prices, costs, results and capital expenditure, and are or may be based on assumptions and estimates related to future technical, economic, market, political, social and other conditions. The Company disclaims any intent or obligation to publicly update any forward looking statements, whether as a result of new information, future events or results or otherwise. The words "believe", "expect", "contracted", "anticipate", "indicate", "contemplate", "target", "plan", "intends", "continue", "budget", "estimate", "may", "will", "schedule", "planned" and similar expressions identify forward looking statements. All forward looking statements contained in this ASX release are qualified by the foregoing cautionary statements. Recipients are cautioned that forward looking statements are not guarantees of future performance and accordingly recipients are cautioned not to put undue reliance on forward looking statements due to the inherent uncertainty therein.*