
Global Testing Proves NOL8's V1.0 Over Multi-Continent Megaport Network, Demonstrating Technical and Economic Superiority Over Legacy Software

Live testing with NOL8 and Megaport traversing the US, Sydney and Tokyo demonstrated agentic AI data workflows at near-zero latency — with global reach to enterprises running Private AI, connecting through Megaport's private network.

Highlights

- **Performance met or exceeded expectations** across a range of conditions, extending the benchmark results released to the ASX on 1 April 2026
- **Live, multi-continent Agentic AI test environment integrated across Megaport links.** Pre-embedding, inference and agent-to-agent simulated use cases run from Virginia, Tokyo and Sydney
- **100% accurate data governance across all use cases and all geographies tested**, with every output verified byte-for-byte
- **NOL8 engine latency measured at 3 milliseconds** in-region, consistent with past tests, with distance the only variable across continents due to the speed of light
- **Every figure reported is a floor, not a limit of the technology.** NOL8's ceiling was never reached; the larger the workload, the wider NOL8's advantage over legacy software
- **10X more governed data per dollar of infrastructure** compared to legacy software on CPUs, a number that will grow as workloads increase beyond the scope of this test
- **Next step: continued global scaling of the NOL8 V1.0 platform with design partners** across the Megaport global network

FortifAI Limited (ASX: FTI) ("FTI" or the "Company") advises the results of the most comprehensive testing program yet completed on the NOL8 AI Data Plane. Performance met or exceeded expectations on every test, and NOL8 was proven across three continents on Megaport's global network for AI use cases. The results also demonstrated NOL8 delivering 10X the amount of governed data per dollar of infrastructure, a powerful advantage in total cost of ownership.

Performance meets or exceeds expectations. NOL8 governed data at the rate of 34 terabytes/day while the legacy software only reached 2.6 terabytes/day, a 10X+ advantage, with every output checked against an independent reference, and its throughput declined less than 1% as governance policies grew by 400%. NOL8's ceiling was never reached, so every figure reported is a floor, not a limit. Matching NOL8 in software would take exponentially more CPUs and their related power draw and still not match NOL8's latency. NOL8 demonstrates robust scalability.

Proven globally, for enterprises running Private AI. Using Megaport's global network — which provides private connectivity to all major global hyperscalers and a global ecosystem of NeoCloud GPU providers and private data centres — NOL8 returned byte-identical governance outputs across Megaport's Virginia, Tokyo and Sydney sites, with distance the only variable. For large enterprises, that means one policy governing AI data wherever workloads run, over private connections, giving them the freedom to move data to any cloud in the world with complete governance control.

Proven cost superiority - more AI per megawatt. Comparing data throughput with hardware cost, NOL8 delivered over 10X more governed data per dollar than legacy software on CPUs. The same testing measured the NOL8 AI Data Plane removing up to 64% of data from the AI pipeline before it reached the GPUs — meaning every megawatt of GPU capacity delivers up to 2.8X more useful AI work, while NOL8 itself consumes a fraction of the power of a CPU-powered software pipeline. As the industry focuses on infrastructure efficiency, this economic advantage will be an important element of customer decision-making.

Next Steps: On the strength of these results, the Company is moving to global scaling of the NOL8 V1.0 platform with design partners across the Megaport global network. Timing and conversion remain subject to partner arrangements.

"This is the applied proof of the NOL8 V1.0 engine. Performance met or exceeded our expectations on every test we ran, and when we took the same workloads to three continents across the Megaport global network, NOL8 returned the same answer, byte for byte, from every one, with a 10X cost advantage. That is what large enterprises running Private AI need: the same cost-effective governance with the freedom to choose between cloud providers, over private connections. That is the foundation for scaling NOL8 globally with our design partners."

— **Kelly Herrell, Chief Executive Officer, FortifAI Limited**

"Enterprises increasingly want to run AI wherever it makes the most sense, across clouds, GPU providers and their own data centres, with their data governed consistently everywhere. As AI infrastructure becomes more distributed, enterprises will need new ways to move data privately between these environments while maintaining consistent governance. NOL8's work reflects the growing industry focus on addressing this challenge, while also highlighting the increasing importance of the private, flexible connectivity provided by the Megaport network."

— **Michael Reid, Chief Executive Officer, Megaport**

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

Media and Public Relations

Matt Wolpin | matt@nol8.io |

Forward-Looking Statements: This announcement may contain forward-looking statements that involve risks and uncertainties. Performance and test figures are drawn from the Company's own testing across defined test corpora and live cloud environments in Virginia (US), Tokyo and Sydney; they are measured results at defined load levels, not maximums, and may not be representative of all production workloads, data types or policies. Data- and token-reduction figures describe the test corpora used and will vary with customer data and policies. Software comparisons from this testing program refer to the Company's own benchmark software implementation, not to any third party's product; references to exponentially growing software shortfalls with policy complexity refer to the Company's benchmark results published to the ASX on 1 April 2026. References to Megaport and other third-party providers do not imply endorsement. References to deployment across hyperscalers, NeoCloud providers and private data centres describe connectivity reach and architectural capability; cross-site testing was conducted on public-cloud environments. Statements regarding design-partner scaling reflect current arrangements only; their timing, expansion and conversion are not assured. Actual results may differ materially from those anticipated. FortifAI undertakes no obligation to update forward-looking statements except as required by law.

About Nol8

NOL8 delivers the AI Data Plane, purpose-built to govern and deliver the right data to every AI agent at wire speed. By enforcing policy inline, NOL8 recovers idle GPU capacity and lifts effective utilisation for NeoClouds and the customers they serve. Learn more at nol8.io.

About FortifAI

FortifAI Limited (ASX: FTI) is an Australian Securities Exchange listed technology company focused on artificial intelligence and advanced computing. In addition to developing AI-forward technologies through its subsidiary Nol8, FortifAI has developed a broad portfolio of video games. FortifAI uses AI to target efficiencies and expansion opportunities in technology.