
Nol8 unveils Roadmap for the Agentic AI Data Plane

Highlights

- Roadmap launched for Nol8's Agentic AI Data Plane Technology.
- Nol8 is in **active conversations with enterprise partners** across AI verticals to design and implement data benchmark testing in high-performance environments, **commencing July 2026** ahead of its **revenue-ready commercial platform launch by end CY2026**.
- Demonstration Engine has proven ability to reduce infrastructure overhead from **5,000 Central Processing Unit's (CPU's) to a single FPGA (Field-Programmable Gate Array)** whilst improving latency from **500 milliseconds (ms) to 3ms (160x improvement)**.
- The AI Data Plane is designed to eliminate the "scalability ceiling" and enable the next generation of Autonomous AI Agents.
- AI Data Plane is the high-speed bridge between Large Language Model inference and execution – Nol8 operates in real-time at the speed of the data stream, known as Data-In-Motion.
- Nol8's current technology has already achieved a leap from **5,000 to 2,000,000 Events Per Second (EPS) (400x improvement)** whilst drastically reducing the space and power required to run mission-critical workflows.
- The Agentic AI market represents a USD\$4T opportunity by 2030ⁱ. This growth will expand global data flows from 334ZB in 2025 to 19,267ZB by 2030ⁱⁱ, driven predominantly by Autonomous AI-generated unstructured data.

Following the completion of its acquisition of Nol8, FortifAI Limited (ASX: FTI) ("**FTI**" or the "**Company**") is pleased to announce the launch of the formal roadmap for its **AI Data Plane**, a revolutionary data infrastructure layer developed by Nol8.

The Nol8 AI Data Plane is designed to eliminate the "scalability ceiling" and enable the next generation of Autonomous AI Agents.

By shifting data processing from traditional software run on CPU clusters to specialized Neural Networks optimized for FPGA-based architecture, Nol8 has already demonstrated the ability to slash infrastructure overhead from **5,000 CPU's to a single FPGA**, while simultaneously improving latency from **500ms to 3ms**.

Introduction to the Artificial Intelligence Data Bottleneck


Artificial intelligence adoption is driving a structural shift in global data infrastructure with the global AI market is forecast to reach multi-trillion dollar scale by 2030, with software revenue already exceeding US\$300 billion annually. As enterprises move beyond early experimentation with large language models such as OpenAI’s GPT, Anthropic’s Claude and Google’s Gemini, deployment is expanding into real-time autonomous systems capable of continuous inference and action.

Traditional large language models operate primarily in prompt-response mode, generating outputs based on user input. The next phase of AI adoption involves agentic systems that operate persistently across enterprise environments, drawing on organisational data, interacting with external tools and executing decisions in real time. These systems require continuous high-throughput data ingestion and deterministic low-latency processing, often measured in milliseconds rather than seconds.

More than 90 per cent of new AI-generated data is unstructured, including text, documents, images, audio and video. Processing unstructured data is materially more compute-intensive than structured database workloads. Legacy data pipelines, built for batch processing and structured datasets, are increasingly unable to support the scale, latency and cost profile required for enterprise-grade agentic AI deployment.

Global Data Forecast

data flows exponential growth



Graph Source: The Holon Data Report Part 4: The exponential shift in data generation and storage capacity to 2040
<https://holoninvestments/the-holon-data-report-part-4-the-exponential-shift-in-data-generation-and-storage-capacity-to-2040/>

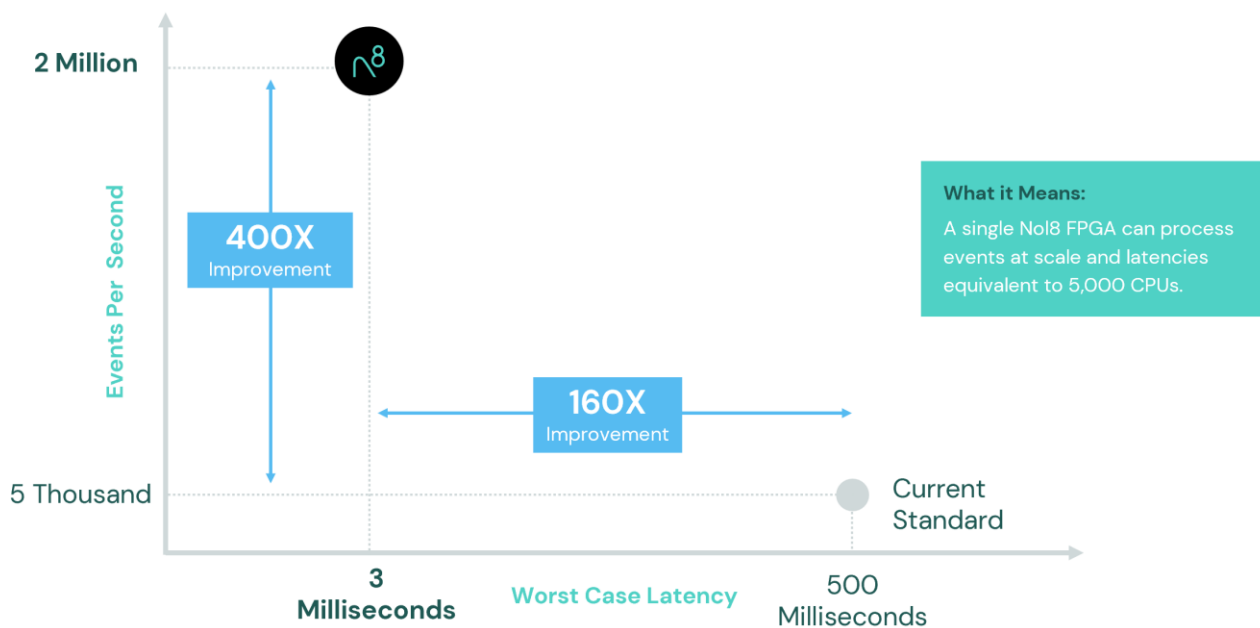
An algorithm built on neural-networks, hyper-accelerated by specialty hardware

This represents an inflection point in AI infrastructure. As AI transitions from isolated model inference to real-time autonomous execution, the performance of the data plane, being the layer that connects models to systems and actions, becomes critical. The limiting factor is no longer model intelligence alone but the efficiency, speed and scalability of data transmission and classification.

Nol8 has developed a neural-network-driven hardware architecture specifically designed to address this constraint. By enabling millisecond-grade processing of high-volume data streams on FPGA-based infrastructure, Nol8’s AI Data Plane provides the performance foundation required for sustained enterprise deployment of real-time agentic AI systems “There is a widening structural gap between having an intelligent AI model and having an actionable outcome,” says Alon Rashelbach, Co-Founder and CTO of Nol8. “True ROI will be determined by the execution layer. The AI Data Plane is the high-speed bridge between large language model inference and execution—an architecture we built to operate at the speed of the data stream itself.”

Proven Performance of Nol8’s Data Streaming Engine

Seismic performance improvements



Performance numbers are estimated using publicly available data from Confluent, Apache Flink, and AWS Flink. Considering 5KB per event, as seen on ClickHouse.

Benchmarking and Commercialisation Roadmap

Nol8 is moving aggressively from proven demonstration scale results to real-world enterprise deployment. The company's roadmap is defined by three distinct stages of engine development:

1. **Demonstration Engine (Proven):** Nol8's current technology has already achieved a leap from **5,000 to 2,000,000 Events Per Second (EPS)**. By drastically reducing the space and power required to run mission-critical AI, Nol8 is providing an environmentally responsible and economically viable foundation for Real-Time Agentic workflows.
2. **Customer Benchmarking Engine (July 2026):** Currently under development by Nol8's engineering team, this engine will enable enterprise to ingest real-world, diverse datasets. Nol8 is currently in conversations with industry partners across verticals to plan, design and implement specific AI data pipeline scenarios in a controlled, high-performance environment.
3. **Commercial Platform (end CY2026):** Following the benchmarking phase, Nol8 will launch its revenue-ready commercial engine. This production-grade platform will be available for first commercial contracts, providing a scalable solution for global organizations requiring real-time intelligence.

Agentic AI a US4T industry by 2030 (McKinsey)

The Agentic AI market represents a four trillion dollar opportunity by 2030, with Fortune 500 companies achieving complete AI adoption. This growth will expand global data flows from 334ZB in 2025 to 19,267ZB by 2030, driven predominantly by Autonomous AI-generated unstructured data types. Traditional infrastructure however, cannot efficiently process this reality, creating critical bottlenecks in latency, compute costs, and the operational complexity enterprises will face.

Market leader Anthropic, forecasting 30x growth in annual recurring revenue to 2030 is already experiencing margin compression from fifty to forty percent due to unresolved infrastructure constraints.

Solving the Scalability Ceiling

The AI Data Plane is designed to decouple data volume from infrastructure cost. By processing and classifying data-in-motion at milisecond-grade speeds, Nol8 allows customers to scale their AI platforms without a linear increase in their "Scalability Ceiling"— the ballooning costs associated with legacy Extract, Transform and Load (ETL) processes and batch analytics.

"The market has moved beyond simple productivity-based AI gains; enterprises now demand AI that can identify and action business processes in real-time to drive revenue," Rashelbach continued. "The introduction of the Benchmarking Engine represents a critical step in dismantling the scalability ceiling. By enabling enterprises to validate high-throughput, millisecond-grade data-in-motion processing against real-world workloads, Nol8 is establishing the execution layer required for scalable Agentic AI. This milestone shifts the focus from model

capability to infrastructure readiness—ensuring AI systems are architected not just for intelligence, but for sustained, future-proofed autonomous action.”

About Nol8 Technology

Nol8 builds 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 environments.

About Fortifai

Fortifai Ltd 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.

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

¹ Superagency in the workplace: Empowering people to unlock AI’s full potential

<https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/superagency-in-the-workplace-empowering-people-to-unlock-ais-full-potential-at-work>