

# Neuromorphic Labs Raises \$5.1M to Build the Trust Layer for Production AI

*Neuromorphic Labs raises \$5.1M to build an enterprise-grade platform that gives teams a faster, more verifiable path from AI experimentation to production.*

SEATTLE, WA, UNITED STATES, August 13, 2026 /EINPresswire.com/ -- [Neuromorphic Labs](#) today

“

Our vision is to make building, deploying, and operating AI as seamless as modern software development.”

*Chetan Kapoor, Founder,  
Neuromorphic Labs*

announced an oversubscribed \$5.1 million seed round led by Flying Fish, with participation from Toyota Ventures, Amplify Partners, Tandem Ventures, SaaS Ventures, and others. The company is building the trust layer for production AI, giving organizations of all sizes a faster path to developing, deploying, and integrating models and agents into business-critical applications and environments, with control and verifiability at every stage.

AI is causing a paradigm shift around application

development and deployment, much like cloud computing did a generation ago. This time the elastic resource is intelligence itself, not compute and storage. Today, building well means finding the right model for each application, optimizing across cost, performance, data locality, and control. Frontier models will carry an enormous share of that work, especially as organizations build proofs-of-concept and validate product-market fit. As AI becomes embedded in production and physical systems, purpose-built models will increasingly become the foundation for differentiated products and services, a shift competitive open-source models have made dramatically more accessible. Across frontier AI labs, AI-native startups, and enterprises of every size, models, weights, data, and agents are becoming core intellectual property. AI is no longer simply a capability organizations consume. It is the strategic infrastructure they own, operate, and increasingly depend on.

Owning AI means owning its entire chain of trust. As AI assets move across developers, pipelines, clouds, production and physical systems, maintaining visibility, control, and trust becomes far harder. Existing tools were never designed for this world, forcing teams to stitch together fragmented infrastructure, security, governance, and compliance solutions to establish provenance, enforce policy, maintain operational visibility, and demonstrate trust across the AI lifecycle. The result is operational complexity, manual review cycles, slower deployments and constrained compute capacity at precisely the moment organizations need to move and scale faster.

Neuromorphic Labs embeds trust directly into the AI lifecycle instead of treating it as a separate review process. Models, weights, datasets, agents, and the compute they run on are given verifiable identity and tracked through every change to enable provenance, consistency, and policy enforcement at run-time. The platform works with the tools AI engineers and data scientists already use and runs inside the environments organizations already operate, across public clouds, on-premises, edge and hybrid infrastructure, while continuously providing the visibility, policy enforcement, and verifiable evidence that business leaders, security, and compliance teams require. Developer productivity and enterprise accountability stop being a tradeoff, because trust becomes an intrinsic property of how AI is built and operated.

"I've had a front-row seat to every major wave of AI infrastructure over the past decade. Every wave expanded what organizations could build, and made the stack underneath exponentially more complex," said Chetan Kapoor, CEO and Founder of Neuromorphic Labs. "As organizations increasingly build and own the models and agents that power their business, trust can no longer be an afterthought. Our vision is to make building, deploying, and operating AI as seamless as modern software development, with trust built into every stage of the lifecycle."

"The next decade of AI will not be defined solely by more capable models. It will be defined by the platforms that let organizations build, deploy, and operate AI systems with confidence, at scale and across diverse compute environments," said Frank Chang, Managing Partner at Flying Fish. "Chetan and his team combine deep market and technical expertise with a clear vision, and we are excited to partner with them as they build this foundational trust layer for the next generation of AI applications."

The funding will expand the founding team and accelerate the platform toward general availability. Before founding Neuromorphic Labs, Kapoor served as Chief Product Officer at CoreWeave, where he helped scale the company through its IPO, and previously led the AI infrastructure business at AWS, growing it 20x into a multi-billion-dollar run rate business. The company is hiring across several roles. To learn more or request early access, visit [www.neuromorphiclabs.ai](http://www.neuromorphiclabs.ai).

Chetan Kapoor  
Neuromorphic Labs

[email us here](#)

Visit us on social media:

[LinkedIn](#)

---

This press release can be viewed online at: <https://www.einpresswire.com/article/933663963>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable

in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.