

Beyond GPUs: Energy Efficient AI

A new learning architecture built for intelligence-per-watt
– from data center to edge

The AI Power Wall


- AI demand is scaling faster than infrastructure
- Energy and cost are becoming hard limits
- Intelligence can't become ubiquitous if compute stays this expensive

Efficiency is no longer an optimization – it's the constraint.



Brute Force Doesn't Scale

- Bigger models demand bigger clusters
- More compute means more power and cooling
- Latency and reliability degrade in real deployment
- The industry needs a new efficiency curve



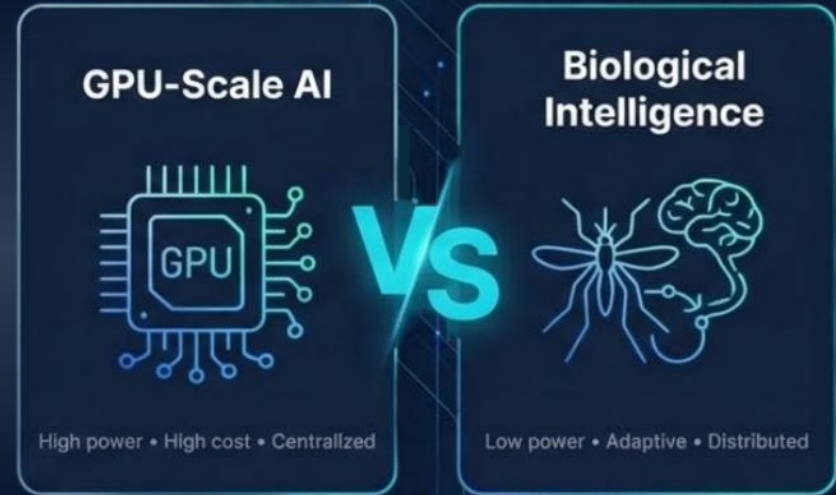
Where the Energy Actually Goes

- Large distributed systems constantly synchronize state
- Token growth increases coordination overhead
- Moving data often costs more than computing

This isn't a tuning problem. It's architectural.

The Efficiency Gap

- Biology achieves intelligence at tiny power
- Modern AI requires massive compute and energy
- This isn't a tuning problem — it's an architecture problem



Why we're the ones asking



Jerry Felix
Chief Architect

 [@jerfelix](#)

AI Visionary & Inventor: 40+ years of experience, multiple patents.

Efficiency Architect: Innovator, reducing training costs, time and energy consumption.

Entrepreneurial Success: Bootstrap founder, 25 years, Successful exit.

Cybersecurity Pioneer: Author of seminal 1987 work on system security and early architect of HP's security consulting offering.



Steve Brunker
CEO

 [@stevebrunker](#)

M&A Technology Strategist: Directed infrastructure M&A for 14 corporate transactions.

Transition Leadership: Guided technology roadmaps at private equity and publicly traded firms.

Enterprise Advisor: 18+ years provided technical strategy to HP's largest global accounts; 20+ years as Enterprise CIO.

The Estimator

- The simplest device that learns a binary pattern
- Builds a usable prediction model in a single pass




Location

- Where we learn: Locally
- Where we infer: Direct local path
- Where we live: At the endpoint

Early → Learning → Converged

Learning Fabric / Inference Mesh

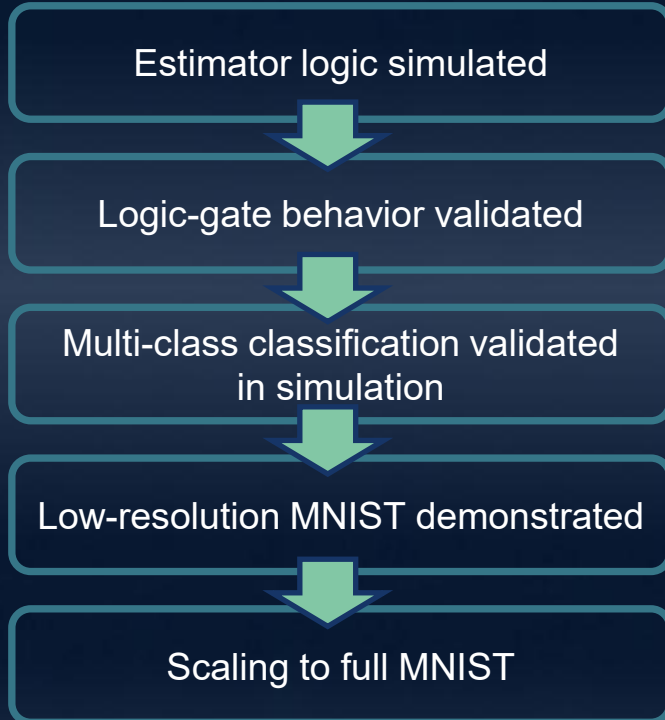
- A uniform tessellated compute fabric
- Local interactions create system-level learning
- Designed for low-energy inference at scale



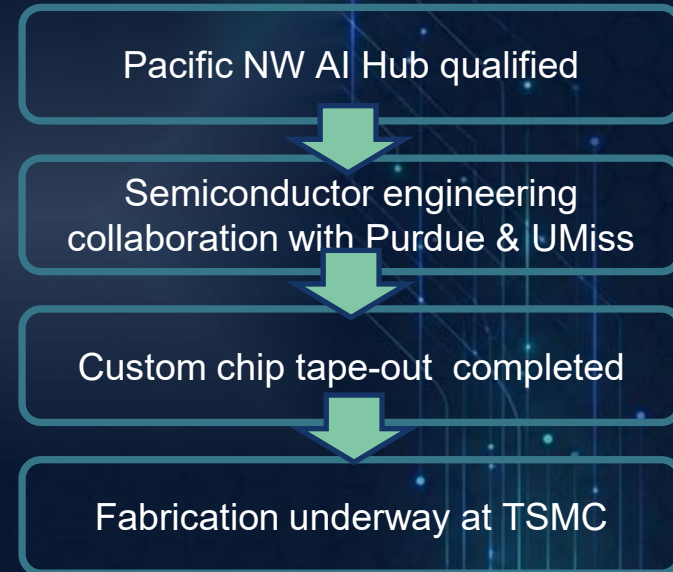
Learning builds structure.
Structure accelerates inference.

From Concept to Silicon

Software Validation



Silicon Execution



Defensible by Design

- Core architecture protected by issued patents
- Covers both the elemental learning device and the scalable fabric
- Multiple filings strengthen long-term platform defensibility
- Built for hardware implementation, not just theory



Estimator

Elemental learning primitive

Learning Fabric / Inference Mesh


Distributed learning architecture

Continuations and International

Broadened patent coverage over time

Technical Validation & Momentum

- Presented research at ISCA (2024, 2025)
- “Most Innovative” award at ISCA 2025 energy-efficiency workshop
- Core architecture protected by issued patents (with additional filings pending)
- Strong collaboration network across academic + industry partners



First Segment: Edge Intelligence

- Real-time classification near the sensor
- Power- and cost-constrained deployments
- Reduced bandwidth and cloud dependency
- Faster decisions, lower total cost



Vision

Cameras + inspection



Audio

Monitoring + detection




Sensor Fusion

LIDAR + industrial sensing

ROI: lower power +
lower inference cost

Go-To-Market Strategy

- Prove performance and efficiency first
- Win a wedge segment with clear ROI
- Expand into larger apps as the platform scales



Our Mission

Make intelligence scalable,
affordable, and sustainable —
by reinventing the architecture
of learning.



Efficiency



Scale



Intelligence Everywhere