

THOMAS BRADY

jthomasbrady@proton.me · thbrdy.dev

Product leader who takes nascent research capabilities and turns them into new product categories. At AWS, worked directly with researchers across graph ML, causality, and LLMs to identify which capabilities had real-world potential — then led 0-to-1 product development from research concept to shipped products (Neptune ML, first LLM strategy, Neuron SDK). Now building full-time as a founder: prototyping daily with Claude Code, shipping an iOS app in 3 days, and publishing original research on AI safety infrastructure. Before tech, 15 years in national security (Special Forces, Senate defense policy) working in high-ambiguity environments where identifying the real structure of a problem is the whole game.

EXPERIENCE

Founder — Pando Industries

2022–Present

- Designed and shipped Notice to TestFlight with beta testers in 3 days using Claude Code—iOS app for interoceptive awareness training with two-tier AI architecture: Apple’s on-device Foundation Models for real-time context, Claude API for longitudinal pattern analysis. First iOS app; zero prior Swift experience. thbrdy.dev/writing/notice
- Published original thesis on claim-dependency tracking for AI safety cases—argues that safety cases are dependency graphs nobody maintains, proposes Scholion as a system for extracting Toulmin-structured arguments from research documents and propagating invalidation when foundational claims are contested. Built working prototype. thbrdy.dev/writing/scholion
- Published technical research identifying 15–30% throughput gap in ML compiler optimization from sequential decision-making across fusion, checkpointing, parallelism, and precision—proposes learned joint planner as alternative to heuristic pipelines. thbrdy.dev/writing/learned-compilation
- Won 1st place at AI in Motion Lab Automation Hackathon (50 builders, 24 hrs)—built multi-agent system with live integration between Openrons liquid handler and Cephla microscope, using LLM orchestration to generate protocols from meeting discussions and close feedback loops with imaging data. [Recap](#)

Senior Director, Product Management — H2O.ai

2021–2022

- Led product strategy for H2O.ai’s core AutoML portfolio representing 85% of revenue, driving 40% customer base growth and \$400K increase in average deal size
- Co-developed H2O Feature Store with AT&T through daily hands-on technical collaboration—built demo notebooks, validated prototypes through iterative customer feedback; attracted \$50M+ pipeline including Workday, PayPal, and Bank of America
- Transformed Driverless AI into cloud-native service, reducing deployment time 10x; ran AutoML experiments alongside CTO, testing features directly with customers

Product Manager — Amazon Web Services

2018–2021

Product Lead, HPC & Elastic Fabric Adapter (2021)

- Authored HPC/EFA strategy positioning AWS in \$10B market; designed pricing strategy balancing \$2B addressable market vs. cannibalization risk; drove wins with Rescale (\$45M), Siemens (\$112M), Rivian, and Boom Aerospace

Product Lead, AI & Machine Learning Research (2020–2021)

- Worked directly with researchers across graph ML, causality, LLMs, and deep learning compilers to identify which research capabilities had product potential; authored research-to-product strategy for \$20M org (200 people)
- Shaped AWS’s first LLM (M5/M*) with Alex Smola, George Karypis, and Trishul Chilimbi—worked in notebooks alongside research teams, defined use cases for Ads/Alexa/Search that grounded research in commercial applications
- Built evaluation framework that doubled research commercialization rate—assessed real-world applicability, technical readiness, and market potential through structured notional press releases, working prototypes, and customer validation
- Launched Neptune ML from research concept to production service, now powering RAG applications; drove distributed training framework delivering 4x performance at 7x cost reduction
- Drove turnaround of AWS deep learning toolchain (now Neuron SDK) through unified approach, reducing inference latency and operational costs by 60%

Product Manager/Solutions Architect, Solutions Architecture (2018–2020)

- Scaled technical workshop program from pilot to 40K+ events generating \$3.4B pipeline; built automated resource provisioning platform enabling 4K+ solutions architects to deliver across 175+ AWS services

Defense Policy Advisor — U.S. Senate Armed Services Committee 2015–2017

- Architected transformation of Army’s \$3B intelligence enterprise—evaluated competing technical approaches, developed new requirements and procurement strategy that cut yearly costs by 65%
- Created investment framework for evaluating emerging defense technologies (IT systems, advanced materials, autonomous systems), guiding multi-billion dollar Congressional investment decisions

Special Forces Communications Sergeant — U.S. Army Special Forces 2004–2012

- Selected from top 5% of 400+ candidates for Special Forces Qualification Course; completed two-year pipeline including advanced communications specialization and intensive Mandarin Chinese immersion
- Engineered tactical and strategic communications architectures for 60+ missions across Asia; led 12-person operations support team managing logistics, comms, and resource allocation for six SF teams

Earlier: Intelligence Analyst, Defense Intelligence Agency (2004) · Manager, Foreign Policy & Research Analyst, U.S. Senate Foreign Relations Committee (1998–2004)

TECHNICAL SKILLS

Python, SQL, Claude Code, LLM orchestration, multi-agent systems, Swift/SwiftUI, data visualization, A/B testing, AWS infrastructure (SageMaker, Neptune, Trainium), AutoML systems, graph algorithms

AWARDS & CERTIFICATIONS

1st Place, AI in Motion Lab Automation Hackathon (2024) · AWS Certified Machine Learning – Specialty (2020) · AWS Certified Solutions Architect – Associate (2019)

EDUCATION

University of Washington — B.A. International Studies (Asian Studies), Mandarin Chinese fluency