

# Comprehensive Briefing: AI in Business

Where AI companies sit, how the workforce actually splits, and which students come out ahead.

AI Vols · April 21, 2026 · University of Tennessee, Knoxville

## Executive Summary

Meeting #5 pivoted from tool comparisons to the business mechanics underneath them. The goal: help students understand where companies compete in AI, what the AI workforce actually looks like, and what career strategy follows from the data. The session closed with a live case study from Revora Systems, the founder's own company.

Key takeaways:

- **B2B SaaS on AI is the most durable business shape today** — model + workflow + proprietary data creates a defensible position that pure wrappers lack.
- **Only 1.4% of AI employees train models.** The rest build products (60%) or operate the business around the models (39%). Career strategy should follow the distribution.
- **Domain knowledge is the new moat.** AI flattens the technical layer, which amplifies the value of industry-specific expertise.

## Types of AI Businesses

AI companies sit on a spectrum — from building the foundation model itself to assembling what already exists. Each position has a different moat profile.

- **Raw LLMs** — GPT-5, Claude Opus, Gemini. A handful of labs, billions in compute. Hardest layer to enter.
- **GPT Wrappers** — thin apps calling an API. Fast to ship, easy to clone, thin moat.
- **B2B SaaS on AI** — Harvey (legal), Cursor (code), Glean (enterprise search). Deep workflow wrapped around the model. Strongest position today.

## What's an API

- **Definition:** the interface by which one piece of software talks to another.
- **Menu metaphor:** your app orders, another service cooks and delivers.
- **The unlock:** you don't need to train a model. Rent one through an API and build the product on top. This is the structural reason a solo founder can now ship what used to require a research team.

## The AI Workforce — By the Numbers

Ravio's 2026 Compensation Trends report covered 400K+ employees across 1,500+ tech companies. The breakdown:

- **60% AI/ML Engineers** — build with existing models, wire APIs into products, ship.
- **39% Everyone else** — PMs, designers, sales, founders, operators. Running the business around the models.
- **1.4% AI/ML Researchers** — actually train frontier models.

**Strategic implication:** the "become an AI researcher" path absorbs almost no one. The "build with AI" path absorbs the vast majority. Career advice should track the distribution.

## Business Strategy — Two Ways In

### Build-First — make a product, then sell it

- **New product** — something the world hasn't seen. OpenAI shipping ChatGPT.
- **Better version** — 10x an existing tool. Cursor making VS Code AI-native.
- **Direct copy** — clone when the market is big enough. Perplexity vs. Google.

### Problem-First — find a real need, then solve it

- **Necessity** — a pain that has to get fixed. Harvey AI for legal review.
- **Improvement** — people can't live without it once they try it. Copilot, Uber, DoorDash.

Every AI business fits on one side or the other. Build-first rewards technical wedge; problem-first rewards domain knowledge. Match the side to what the founder already has.

## Case Study — Revora Systems

The founder's own service business, targeting industries that modern software bypassed. Five lessons surfaced:

- **Problem-first wedge.** Real workflow pain in under-served industries.
- **Domain edge.** Prior construction-industry experience beats an outsider tech generalist.
- **Distributed execution.** Sales, delivery, support, ops — no single person does it all.
- **AI as one tool.** Sometimes the right answer is a spreadsheet, a phone call, or a trained VA.
- **Compounding leverage.** 3x delivery speed enables 2x price while still out-delivering competitors.

## The Build Process, Then vs Now

- **Then:** CEO → Designer → 2–3 Junior Devs → Lead Dev → Product. Months, full team.
- **Now:** CEO → Lead Dev + AI → Product. Weeks, one driver.
- **Replaced:** designer (AI generates UI), junior devs (AI writes boilerplate), QA grunt work (AI reviews its own output).

- **More valuable:** lead developer, COO/operator, founder. The people left carry more weight.

## Strategic Takeaways

- **Be the top 20%.** Deep domain knowledge + fluent AI use beats either one alone.
- **Skip the middle.** The disappearing roles are the ones students used to aim for as "first jobs." Aim higher or aim sideways — operator, product, domain lead.
- **Build public evidence.** A portfolio of shipped AI-assisted work matters more than any certificate in a hiring market that's already saturated with credentials.

*Horses became cars. The jobs changed. The people who adapted came out ahead.*