BUILDING AI-NATIVE PROFESSIONAL SERVICES FIRMS

Strategy, Economics, and Execution

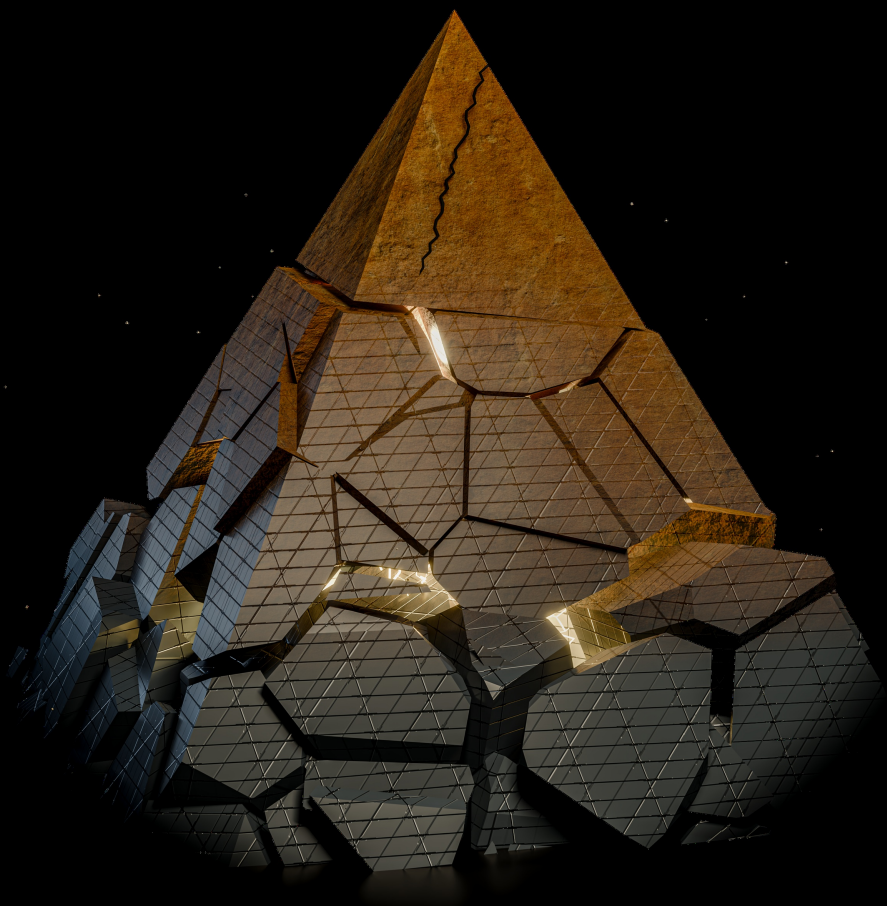

DANIEL M. KATZ · MICHAEL J. BOMMARITO II · JILLIAN BOMMARITO

Chapter 1

The AI-Native Professional Services Firm

Chapter 1 · March 2026

Chapter 1: The AI-Native Professional Services Firm

Through all these misfortunes, these dangerous times,
we head for Latium, where the fates hold peaceful lives
for us: there Troy's kingdom can rise again.

— Virgil, The Aeneid, Book I

Every startup founder has an origin story. A moment when they see what others miss. An experience that breaks the spell of how things have always been done. For some, it arrives as a flash of insight—a technical breakthrough, a market inefficiency too obvious to ignore. For others, it accumulates slowly—a thousand small frustrations that crystallize into a single realization: ***there has to be a better way***.

These origin stories matter because they reveal why founders take the risk. Why they leave stable careers, forgo partnership tracks, and bet everything on an uncertain future. The best origin stories are not about seeing opportunity in the abstract. They are about living inside a broken system long enough to understand exactly what needs to change—and having the conviction to build the alternative.

Sarah Okonkwo's origin story began in a conference room. Not with a dramatic exit or a bold declaration, but with a presentation that revealed the gap

between what her firm claimed to be doing with AI and what AI actually meant for the future of professional services. That gap would become the foundation of everything she would later build.

Sarah Okonkwo sat in the conference room on the fourteenth floor of a glass tower in downtown Phoenix, watching her managing partner describe the firm's new AI initiative. It was the second week of July 2024—the firm had returned from the holiday weekend to this all-hands meeting—and Sarah had walked in genuinely hopeful.

"We are going to license Microsoft Copilot," the managing partner explained, clicking to a slide showing a generic interface. "This cutting-edge AI platform will help our associates accelerate contract review, draft routine correspondence, summarize depositions, and analyze regulatory compliance matters. We will be able to handle larger matters more effectively and efficiently."

Sarah glanced around the table. The other partners nodded approvingly. One asked about billable hour credit for supporting tool usage. Another wondered if they should disclose AI use to clients. No one asked the question that kept Sarah awake at night: *If AI capabilities are advancing this rapidly, what happens to our business model in the years to come?*

Sarah had been a computer science minor in college—an unusual pairing with her pre-law track, but she had always been interested in technology. She understood what the managing partner did not: Copilot was fine, but it was not the point. The firm was licensing a consumer tool and calling it innovation. They were playing at the margins when there appeared to be a much more fundamental shift occurring in the delivery of knowledge work.

She had joined the firm eight years earlier, straight from law school. She had made senior associate in five years, putting her on the partnership track. She billed 2,200 hours annually, mentored junior associates, brought in clients. By every traditional metric, Sarah was succeeding.

Her mid-market firm billed her time at $675 per hour (rack rate). Her fully loaded cost—salary, benefits, overhead—was roughly $300,000 annually. Sim-

ple arithmetic: she generated well over a million dollars in billings and cost the firm a fraction of that. The surplus funded partner compensation. That was the deal. Work hard, bill hours, make partner, collect your share of the surplus the next generation created.

Except AI was about to put significant pressure on that historic arrangement.

Soon after GPT passed the bar exam in March 2023, Sarah had started experimenting with LLMs in her spare time. Not just Microsoft Copilot—her tinkering had taken her deeper. By mid to late 2023 she had already been testing Claude, GPT-4, Palm, and even Llama directly, learning how to prompt them, understanding their failure modes, exploring what they could actually do versus what vendors claimed. She had stayed in touch with several computer science friends from undergrad, some of whom now worked at frontier model companies. Their late-night conversations about model architectures, training approaches, and emerging capabilities gave her insights the managing partner would struggle to acquire.

That previous fall, the firm had announced an AI strategy committee—a small group of senior partners and the IT director, tasked with evaluating how the firm should respond to the technology. Sarah had hoped to be included. She had mentioned her experimentation to the managing partner, offered to share what she had learned. The managing partner thanked her and said the committee would reach out if they needed input from the associate level. They never did. Sarah was not offended—she understood the firm's hierarchies—but she followed the committee's work from a distance, curious about what they would recommend. When the all-hands meeting appeared on the calendar for the week after the Fourth of July, she assumed the committee was ready to present its findings. She allowed herself to think the firm might surprise her.

She was quite familiar with Microsoft Copilot. She had tested it extensively. For simple tasks—drafting routine emails, basic summarization—it was fine. But for even modestly complex legal work, the kind that required deep contextual understanding, nuanced analysis, and the ability to spot subtlety in large amounts of information, Copilot had severe shortcomings. It could not hold enough context. It could not reason through edge cases. It could not learn from

the firm's institutional knowledge because it had no institutional knowledge to learn from. Yet this was what her firm was betting on.

Despite the shortcomings of Copilot on real legal work, the results from her direct experimentation shocked her. The initial draft of a contract analysis that would have taken her eight hours took an LLM a matter of mere minutes. Yes, she spent three hours reviewing, correcting, and refining. But the total time collapsed from eight hours to roughly four hours.

That is why the managing partner's announcement landed the way it did. Eight months of committee work, and the recommendation was to license Copilot—a decent tool for what it was, a general product for the average business professional. Sarah had allowed herself to hope for more. The real opportunity required building something custom, something that compounded knowledge, something that fundamentally changed how work flowed through the firm. The committee had not seen it. Or perhaps they had seen it and decided it was too ambitious, too risky, too far from what the partnership would accept.

The next time she was presented with the same type of task she had tested previously in her experiment with an LLM, she did not build upon her experimentation. Instead, she did the work the traditional way and billed the client for eight hours. But she wondered: *When the firm figures this out, what happens to the other four hours?* More importantly: *When clients figure this out, why would they pay $675 per hour for the elements of the work that AI can do in minutes?*

The managing partner's presentation continued. "We see this as a competitive advantage. We can deliver higher quality work in the same or less time, which our clients will value."

Sarah raised her hand. "Are we changing our pricing model?"

Silence.

"I mean," she continued, "if we are completing work in less time—perhaps even in say half the time—are we charging half as much? Or are we billing the same amount for less work?"

The managing partner smiled—the smile partners give associates who ask inconvenient questions. "Let us take this offline but to respond quickly—we

will continue to bill based on the value we deliver. If we are more efficient, that benefits our clients through faster turnaround."

Translation: I am not going to address the elephant in the room even though it is not going anywhere.

Sarah walked back to her office quietly disappointed. She did not blame the committee—they had done what the partnership asked them to do. But the opportunity was right in front of them. A once-in-a-generation shift in how professional services could be delivered. AI was not just another efficiency tool like a document management system or a KM platform. In her mind, it was a fundamental restructuring of the production function. And the firm's response, after eight months of study, was to license Copilot and call it a strategy.

Her computer science background told her what was coming. The technology had already crossed some sort of a threshold. By July 2024, large language models could reason about complex text, draft professional-quality outputs, and learn from feedback. They weren't perfect—they made mistakes, missed context, occasionally hallucinated—but the leading frontier models were good enough to handle first-pass work on a number of tasks that junior associates did. And with agentic frameworks, tool calling, and a range of other techniques, they were only going to improve even more rapidly.

The firm could have been building proprietary systems, capturing institutional knowledge, redesigning workflows from the ground up. Instead, they were moving too slow, unwilling to cannibalize their own business model, hoping the disruption would happen gradually enough to manage with incremental adjustments.

Sarah ran calculations. The professional services industry—law, consulting, accounting, financial advisory—generates several trillion dollars annually worldwide. Legal services alone exceeds one trillion. This is one of the largest and most profitable sectors in the global economy, operating on a business model unchanged for a century.

AI threatened to put that model under significant pressure. And her firm was responding by buying a general purpose business tool and basically pretending nothing fundamental had changed.

That night, Sarah made two decisions. First, she would go deeper—exploring the leading frontier models, testing agentic capabilities, experimenting with tool calling and multi-step reasoning to truly determine what the future operating model might look like. Not AI as a tool bolted onto existing workflows, but as the foundation for a completely different way of delivering professional services. Second, if her firm would not build that model, she would find someone who would. Or perhaps she would need to build the future all by herself.

1.1 The Trillion-Dollar Industry

Within a few months of that conference room presentation, Sarah would no longer be at the firm. But to understand why she left—and what she built instead—requires understanding the economics she was trying to escape.

Professional services firms sell expertise typically delimited in units of time. Clients pay for judgment, analysis, and advice they cannot easily produce themselves. The work is knowledge-intensive, relationship-driven, and resistant to commoditization—or so the industry has long believed.

1.1.1 *Market Size and Structure*

The global professional services market spans four major sectors, each generating substantial revenue and employing distinct business models.

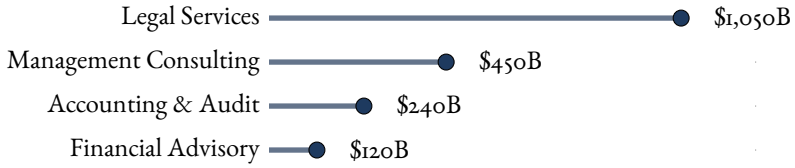

Figure 1.1: Global Professional Services Market by Sector, 2024. "Accounting & Audit" represents auditing services only. See Appendix **??** for sources and methodology.

Each sector has distinct market leaders: the Am Law 100 (generating $160 billion of the legal total), MBB and Big Four advisory practices in consulting, Big Four

networks in accounting, and investment banks with specialized wealth managers in financial advisory.

When Sarah studied these numbers in the spring and early summer of 2024—while the committee deliberated behind closed doors—she saw something her former partners did not: vulnerability masquerading as strength. These markets share common characteristics. They are fragmented, with thousands of firms ranging from solo practitioners to global networks. They are profitable, with partner compensation reaching millions at elite firms. They are growing, driven by regulatory complexity, globalization, and demand for specialized knowledge.

These strengths, however, mask structural vulnerabilities that Sarah had experienced firsthand.

1.1.2 *The Staffing Pyramid Model*

Professional services firms create value through the staffing pyramid. A partner supervises a team of associates, each billing at rates exceeding their fully loaded cost. The pyramid structure—many associates per partner, many junior staff per senior—generates the margin funding partner compensation.

Sarah had been part of that pyramid. Her practice group ran typical numbers: one partner supervising six associates. The partner billed at more than $1,000 per hour, she billed at $675, junior associates starting at $395. The partner worked maybe 1,400 billable hours annually—client development, firm management, and mentoring consumed the rest. Sarah and her peers billed 2,000–2,400 hours each.

Simple math: the partner generated more than $1.4 million in billings. The six associates generated $5.83 million combined. Total maximal practice group billings: $7.23 million (not counting utilization and realization / markdowns). Total compensation for the group: roughly $1.9 million for the associates (Sarah and her peers made $220,000–300,000 plus certain bonuses).

But that number was misleading. The practice group did not exist in isolation. The firm had to pay for office space, support staff, technology, insurance, marketing, and administration. Those costs consumed another 30 to 40 percent

of revenue. And what remained after overhead did not accumulate as retained earnings—partners distributed virtually all of it each year as compensation. By the time the firm paid its obligations and its partners took their draws, there was little left to reinvest. For an outside investor looking at the business, the actual operating margin—the money available after all costs including partner compensation—was thin, often in the range of 10 to 15 percent (or less). The practice group economics looked impressive. The enterprise economics told a different story.

Make no mistake, the model worked beautifully—for the partner.

The Am Law 100 illustrates this model at scale (**amlaw2024**). In 2024, these top law firms generated $160 billion in revenue with 124,000 lawyers. Revenue per lawyer averaged $1.28 million. Profits per equity partner averaged $3.15 million. The spread between these figures represents the value extracted from the staffing pyramid: senior professionals capture the surplus generated by junior professionals.

The model works until it doesn't. Associate salaries have risen faster than billing rates. Realization—the percentage of billed hours clients actually pay—has declined as clients push back on fees. Utilization—billable hours as a percentage of total available hours—hovers around 69 percent across the broader professional services sector (**spi2024benchmark**), below the 75 percent target most firms need for healthy economics. Staffing ratios are shifting: non-equity partners now outpace equity partners in many firms, a sign that at a number of firms the traditional pyramid is already starting to flatten.

Sarah had watched these pressures build. She had seen associates leave for in-house roles or tech companies. She had seen clients reject bills, demanding write-downs. She had seen clients insource more work, outsource work to alternative providers (ALSPs, LPOs, etc). She had seen partners panic about utilization rates while simultaneously blocking investments in technology that might improve them.

But what crystallized her thinking was simpler: she had seen what happened when she used AI to do her work. The pyramid did not just compress—in her mind—it crumbled.

1.1.3 *Value at Stake*

Despite these pressures, the trillion-dollar market generates substantial profits. The Am Law 100 reported $69 billion in net income in 2024 (**amlaw2025**). The Big Four accounting networks generate $38–70 billion each in revenue globally (Deloitte leads at $67 billion (**deloitte2024annual**), KPMG trails at $38 billion (**kpmg2024annual**)). McKinsey, Bain, and BCG command premium rates and recruit from the same talent pools as the most selective technology companies.

But the value is unevenly distributed. Elite firms in each sector capture disproportionate shares of profit. Kirkland & Ellis alone generated $8.8 billion in revenue in 2024 (**amlaw2025**). Wachtell reported a 78 percent profit margin—extraordinary by any measure (**amlaw2025**). At the other end of the spectrum, hundreds to thousands of mid-market firms compete on price, struggle with retention, and watch their margins compress year after year.

Sarah's former firm was solidly mid-market. Not elite, not struggling, just—steady. 120 lawyers, $88 million in revenue, respectable profits per partner. The kind of firm that might otherwise exist in twenty years doing exactly what it did today, serving the same clients, billing the same way, until something disrupted it to its core. In Sarah's estimation, AI threatened to be that disruption.

Sarah had a degree of conviction in her emerging thesis but her assessment was not guaranteed to be correct. Law firms had weathered many storms including wars, financial crises and pandemics and yet the economic model was still largely in the spirit of the system that Paul Cravath had devised more than a century before (**oller2019whiteshoe**).

But Sarah's confidence was growing. She increasingly believed that firms that adopted AI effectively would lower costs, improve quality, and capture market share. Firms that did not would find themselves squeezed from above by elite players and from below by AI-native entrants operating on different economics.

Sarah was working up the confidence to be one of those entrants. But it was a bold leap—a cosmic bet. Leaving a lucrative job with a clear path to partnership was a non-trivial decision. She was getting close to the water's edge. But to fully convince herself—she needed to convince herself as to why this transformation was going to be different.

1.1.4 *Why This Transformation Is Happening Now*

Of course, the claim that AI will transform professional services invites skepticism—Sarah heard it constantly in the weeks before and after she left her firm. Legal technology had promised disruption for decades without truly delivering it. Knowledge management systems, document automation, and practice management software all claimed to transform professional services. Most innovations failed to fundamentally change how work gets done. The most prominent cautionary examples were still fresh. Clearspire, launched in 2011 with a visionary two-entity model separating legal practice from business operations, had shut down by 2014—its founder later acknowledged they had become "slaves to the elegance of our vision" in a market that was not ready for what they were selling (**cohen2017clearspire**). Atrium, founded by Twitch co-founder Justin Kan in 2017 with $75 million in venture backing from Andreessen Horowitz, had collapsed by 2020 (**constine2020atrium**). Atrium's post-mortem was instructive: the concept was sound—use technology to make legal services radically more efficient—but the underlying AI was not yet capable of meaningful legal reasoning, the ecosystem of tools did not exist, and the venture capital created growth expectations that a services business could never satisfy. Kan's own lesson was blunt: "Adding more money to a situation of lack of product-market fit rarely works" (**kan2021pmf**). Both firms had elements of the thesis right. But timing is everything.

Sarah's former managing partner had made exactly this argument when she gave notice. "We have seen technology waves before. We adapted. We will adapt to this one too."

Sarah understood the skepticism. She shared some of it. Yet three forces kept surfacing in her research, forces she could not dismiss:

Technology capability crossed a threshold. Starting in 2023 and in the years thereafter, large language models particularly when coupled with advanced agentic capabilities and knowledgeable human users achieved something previous AI could not: they can read, reason about, and generate professional-quality text. This capability emerged from scale—billions of parameters trained on trillions of tokens of text, essentially ingesting the written output of human civilization.

Earlier AI required narrow, structured inputs—think keyword searches, decision trees, rigid templates. Those systems could only handle what they were explicitly programmed to recognize. Current models learned patterns from such vast quantities of examples that they developed emergent abilities: they handle ambiguity, recognize context across documents, and produce drafts that read like a junior professional wrote them, not a machine.

Although it is ultimately a question of degree, this capability shift is likely to rewrite the economics of professional work. Tasks that once required hours of associate time—the first-pass review, the initial draft, the background research—can now be reduced to minutes. The gap between AI-assisted draft and final work product had narrowed significantly. What remains is the judgment layer: the senior review, the client conversation, the strategic recommendation. The pyramid that professional services firms have built their economics on—leveraging junior talent across large volumes of routine work—suddenly has a different foundation.

Sarah had tested this personally. She had given an agentic system a gnarly indemnification clause from a vendor contract and asked it to identify risks. It found issues she would have found—and two she might have missed on a first read. It was not perfect. It occasionally mischaracterized context or missed subtle connections. But it was good enough to handle the first pass on the vast majority of routine work. That "good enough" is arguably what changes everything. And the tools were only going to get better.

Economic pressure intensified. Associate salaries at major law firms have risen from $160,000 to $225,000 in five years, while billing rate increases have lagged. Partner compensation expectations remain fixed. At her old firm, the margin squeeze was real and accelerating. Bet the company or high stakes scenarios notwithstanding, most firms (like her mid-market firm) were struggling to grow profits simply by raising rates—they had to do something different.

Sarah knew the numbers cold. Her former firm's revenue per lawyer had been flat for three years. Operating expenses were up 13 percent. Maintenance of compensation per equity partner had been only achieved through a reduced real estate footprint, targeted de-equitization, the addition of so called partner tiers among other things. The only variable they could control was staff and

associate headcount—and they had already cut staff to the bone. The next cut would come from associates. Which meant the pyramid was already flattened whether or not they acknowledged it.

Regulatory barriers began falling. Arizona authorized full non-lawyer ownership of law firms in 2020. The UK has permitted Alternative Business Structures since 2012. Texas blessed MSO structures in 2025. When constructed carefully, these openings allow external capital to fund transformation that partnerships cannot finance themselves. In accounting, consulting, and financial advisory, fewer restrictions exist, allowing faster movement.

This was the piece that made Sarah's decision possible. If structured properly, she could build a different type of enterprise, without lockstep compensation, without the kind of partnership politics that had killed every innovation proposal she had watched get debated and deferred at her old firm. She could raise venture capital, hire differently, price differently, and operate on completely different economics.

These forces were mutually reinforcing. Regulatory openings enabled capital formation. Capital funded technology development. Technology enabled productivity gains. Productivity gains attracted more capital. The cycle had begun, and it favored firms that moved early.

Sarah moved early. But, of course, Sarah was not the only individual to see the opportunity for a different kind of legal service organization. In the months and years that followed, hundreds of millions of dollars in venture capital would ultimately be committed to the idea of an AI-Native Firm.

Within months of the presentation about the virtues of Microsoft Copilot, she gave notice and brought one associate with her—Elena, who had been her closest collaborator for three years. Her first hire after that would be Joshua Thornton, a regulatory specialist she recruited from a well-regarded Phoenix firm, whose expertise she would need from day one.

The name for her new entity came on a Tuesday night in September, as she was preparing paperwork. Sarah was at her kitchen table with the entity formation paperwork spread around her laptop, a half-eaten bowl of pasta pushed to the side. The form asked for the name of the proposed firm. She had been cycling through options for days—combinations of initials, geographic references,

Latin words that sounded serious. None of them fit. They all sounded like the firm she was leaving.

She thought about what had driven her out. Not the economics, though the economics were part of it. Not the technology gap, though the gap was real. It was the lack of real honesty about the moment. An AI committee that met for eight months and produced a memo recommending further study. The partners who told clients they were "embracing innovation" while privately dismissing it as a fad. The gap between what the firm said and what the firm did.

She pushed the laptop aside, pulled a yellow legal pad toward her, and started doodling. She wrote names in the margins, crossed them out, wrote more. *Apex*—too aggressive. *Clarion*—too precious. *Vanguard*—taken by an index fund. She drew boxes around the better ones and arrows between the worse ones and filled the margins with small geometric shapes while her mind worked.

Then she wrote a word near the bottom of the page, almost as an afterthought: *Candor*.

She sat with it for a moment. It was not clever. It was not Latin. It would not impress the partners at her old firm, which was part of why she liked it. It said exactly one thing: this firm will tell you the truth—about what AI can do, about what it cannot, about what your work is worth, about what the future looks like. No committees. No memos recommending further study. No Copilot dressed up as transformation.

She circled the word twice, then underneath it sketched a rough logo—two overlapping squares, one dark and one light, the kind of thing you draw without thinking. Transparency. Overlap. Seeing through to what is real. She tore the page off the pad, set it against the laptop screen, and looked at it from across the table.

What would become her AI-native practice—was not yet a firm. It was a name on a government form and a conviction that the industry needed something it was not getting: honesty about what was coming. But at the time she filed the paperwork, Sarah was mostly following instinct. She had a general vision of what a future firm might look like but had not yet fully crystallized her thinking about what exactly constitutes an "AI-Native" Professional Services Firm.

1.2 What AI-Native Means

Understanding why the transformation is happening leads to the next question: what does transformation actually mean? The term "AI-native" is frequently misused. Sarah encountered this in her first conversations with potential investors and clients.

"So you use Microsoft Copilot?" one investor asked.

"We use enterprise AI platforms, yes," Sarah replied. "But that is not what makes us AI-native."

"What does?"

That question forced Sarah to articulate a definition that would become her firm's foundation.

1.2.1 *The Definition*

AI-Native Professional Services Firm

An AI-native professional services firm is one where AI is embedded in the core production system—not as a tool that professionals use occasionally, but as the foundation of how work gets done.

This definition has four operational implications that distinguish AI-native firms from those that merely use AI tools. Sarah learned these implications by building them into her firm's operating model.

Work flows through AI first. Agents complete initial analysis, draft documents, identify risks, and prepare outputs that humans review, refine, and deliver. The sequence reverses the traditional model: instead of humans doing work with AI assistance, AI does work with human supervision.

When an AI Native law firm such as Candor receives a task such as a contract review engagement, the workflow starts with ingestion: documents flow into a system that routes them to specialized AI agents. One agent handles risk classification. Another extracts key terms. A third compares provisions against the client's standard positions. The outputs queue for human review, but the AI consumes a significant percent of the overall task.

At her old firm, to the extent that a client had not already outsourced or insourced the work, an associate would have read every contract, manually. The associate might use an AI tool to help draft summaries, but the human did the work with AI providing support. The sequence determined everything about economics and scalability.

Humans supervise rather than execute. The role of the professional shifts from doing work to verifying, improving, and adding judgment to AI-generated output. This is not a minor adjustment but a fundamental redefinition of what professionals do.

Sarah would ultimately hire for different skills than her old firm valued. She needed people who could read AI-generated contract summaries and spot what the AI missed—the subtle cross-references, the missing carve-outs, the context that required human judgment. She did not need people who could read two hundred contracts in three weeks. She needed people who could verify that AI-reviewed contracts were correct in three days using a scientifically oriented methodology.

Knowledge compounds in systems. Each engagement improves the firm's capability by adding to data sets, refining models, and building institutional memory that does not walk out the door when employees leave. The firm's knowledge base becomes a durable competitive asset.

This was the piece Sarah's old firm could not replicate. At a traditional firm, knowledge lived in partners' heads. When partners retired, the knowledge left. Sarah's firm captured every engagement: client preferences, risk patterns,

successful negotiation positions, edge cases where AI failed. That knowledge fed back into the system, making the next engagement better. The firm's capability compounded.

Economics are technology-driven. Margin comes from the efficiency of the AI system rather than primarily from labor arbitrage. The traditional staffing pyramid economics give way to platform economics with different scaling characteristics.

In the competition with Marcus's firm, Candor spent $8,000 in labor and compute costs to review two hundred contracts—work that would have cost a traditional firm $47,000 to deliver. She charged $22,000, making 65 percent margin while saving the client $25,000. The margin came from technology efficiency, not from paying associates $250,000 and billing them at $675 per hour.

The distinction is important: a firm can use AI extensively and still be fundamentally traditional if the work still flows through humans first, with AI providing support at the margins.

1.2.2 *The Spectrum*

The distinction between AI-native and traditional is not binary. Firms exist on a spectrum from traditional to AI-native. Sarah learned this by watching firms at different stages respond to the same technology.

AI-Enhanced firms use AI tools to support existing workflows. A lawyer might use a large language model to help draft a brief or summarize a document. An accountant might use AI to flag anomalies in transaction data. The underlying work process remains unchanged—humans do the work, AI assists where convenient.

Sarah's old firm was at best AI-enhanced. They had licensed tools. They had trained associates on prompt engineering. They had implemented guardrails and disclosure policies. But the work still flowed the same way it always had: associates billed hours, partners reviewed work, clients paid for time. The economics were unchanged.

Most professional services firms today are AI-enhanced, or aspire to be. They have experimented with generative AI, perhaps implemented some practice-specific tools, and trained their professionals on "prompt engineering" or similar skills. The economics are unchanged: the firm still bills by the hour, still relies on the staffing pyramid, still operates the traditional structure.

AI-Augmented firms have redesigned workflows to incorporate AI at specific stages. A due diligence team might use AI agents to complete a first-pass review of contracts, with lawyers reviewing the AI output rather than reading every document from scratch. A consulting engagement might use AI to generate initial analysis that consultants then refine and present.

Sarah saw these firms as competitors—sophisticated firms that understood AI could improve efficiency but had not yet rebuilt their core operating model. They captured real efficiency gains. They completed work faster, often at lower cost, and could handle larger volumes than traditional firms. But the core economics remained labor-based. Professionals still did substantial work. Billing structures might shift toward fixed fees, but the cost base was still dominated by human labor.

AI-Native firms build their production system around AI from the start. Work flows through AI agents first. Humans supervise, add judgment, and handle the exceptions that AI cannot. The firm's competitive advantage comes from its AI systems, its data, and its ability to compound knowledge—not primarily from the staffing pyramid.

This was what Sarah was building. She was pursuing fundamentally different economics. She could offer fixed-fee arrangements because her costs were predictable. She could handle volume that would overwhelm a traditional firm. She could serve clients profitably at price points that traditional competitors could not match. And her margins would expand as her AI systems improved.

1.2.3 *Characteristics of AI-Native Firms*

Beyond the spectrum position, AI-native firms share several distinguishing features. Sarah discovered these characteristics by building them into her firm—and by watching other AI-native entrants make the same choices.

Technology functions as a core asset. Sarah's AI platform was not a purchased tool but a proprietary system that embodied her firm's knowledge, methods, and competitive advantage. She used frontier models—Claude, GPT, and others—but the integration, workflows, and accumulated knowledge were proprietary. This platform represented the firm's competitive moat, and she invested heavily in building it. The distinction matters: any firm can license the same foundation models, but the proprietary layer—the workflows, the domain-specific training data, the feedback loops—is what creates defensible advantage.

Knowledge capture becomes strategy. Every engagement added to Sarah's institutional memory. Client preferences, deal patterns, regulatory interpretations, and practice innovations were captured in systems that made the next engagement better. When a client requested changes to standard contract language, those preferences were stored and automatically applied to future contracts for that client. When an AI agent missed a risk, the correction fed back into the quality control process. The data moat widened with each completed matter. Traditional firms could not replicate this because their knowledge lived in partners' heads and left when partners retired.

Supervision replaces execution as the core human role. Sarah's lawyers spent less time doing and more time verifying, judging, and improving. The skill set was different. It was less the mechanics of drafting and more about understanding AI outputs, identifying errors, and applying contextual judgment. Sarah hired one senior lawyer who had spent fifteen years at Big Law doing M&A diligence. He was skeptical at first—"I do not know how to supervise AI"—but learned quickly. Within months, he could verify AI-reviewed contracts faster and more accurately than junior associates could read them cold. He had the dexterity to work differently. Most did not.

Fixed or value-based pricing becomes standard. Sarah's firm priced by the project, outcome, or value delivered rather than by the hour. This aligned incentives: the firm profited by becoming more efficient, not by taking longer. When she quoted $22,000 for the vendor contract review that Marcus Chen quoted at $47,000, she was not underpricing to win business—she was pricing based on her actual costs plus healthy margin. As her systems improved, her costs fell while her pricing remained stable, expanding margins.

Capital intensity increases. Building and maintaining AI systems requires investment that traditional partnership economics struggle to support. Partners distribute profits annually; AI platforms require sustained, multi-year capital deployment. Sarah used her early funding to build technology, hire specialized talent, and fund operations during the growth phase when revenue lagged investment. AI-native firms often raise external capital—venture capital, private equity, or strategic investment—because the traditional partnership model cannot fund the necessary technology investment on its own.

1.3 Questions for Sarah, Questions for All of Us

Sarah would spend many months building. It took a while but eventually, she started getting calls—from former colleagues asking how she did it, from other firms asking for advice, from investors asking for introductions. Everyone wanted to understand: how did you see this opportunity? How did you make the jump? How is it working?

This book exists to answer those questions—not just for Sarah, but for everyone facing similar decisions.

1.3.1 *Professional Services Leaders*

If you are a partner, principal, or senior executive at a professional services firm, you face a strategic choice. Your firm can attempt to transform into an AI-native model, capturing the efficiency gains and margin expansion that transformation promises. Your firm can continue on its traditional path, optimizing at the margins while hoping the disruption happens slowly enough to manage. Or your firm can combine with others—through merger, acquisition, or partnership—to build scale and capability that neither could achieve alone.

Each path has risks. Transformation is expensive and uncertain. Standing still leaves your firm vulnerable to competitors who move faster. Combinations

fail more often than they succeed. This book provides various perspectives for evaluating these options and making decisions that fit your firm's circumstances.

Sarah's path—leaving to build something new—is one option. But transformation within existing firms is possible too. Incumbents actually have many advantages not the least of which is a set of existing client trust relationships. Sarah's story as well as the frameworks in this book apply whether you're building from scratch like Sarah or transforming an established practice.

You do not need to be at an elite firm to benefit. The transformation affects firms of all sizes. Mid-market firms may have more flexibility to change than their larger, more complex competitors. Boutiques may find advantages in specialization and agility. Sarah's firm started with three people and few clients. Your transformation may look different, but the principles apply.

1.3.2 *Financial Sponsors*

If you are an investor—private equity, venture capital, family office, or strategic buyer—professional services represents both an opportunity and a puzzle.

The opportunity is clear. A trillion-dollar industry built on an outdated model is ripe for disruption. Firms that successfully transform will capture extraordinary value. In 2025, venture capital flowed into AI-native law firms at unprecedented rates: multiple firms raised rounds in the tens of millions, with valuations in the hundreds of millions, and private equity deployed significant capital through technology licensing structures.

Sarah's firm would ultimately attract investment because investors saw what she saw: a structural opportunity to rebuild professional services on AI-native economics. Her pitch deck first showed traditional firm economics: modest net margins, linear scaling, labor-constrained growth. The second showed her firm's economics: higher margins expanding further as systems improved, technology-enabled scaling, capital-efficient growth. Investors understood the potential.

The puzzle is how to evaluate these investments. Traditional valuation approaches—revenue multiples, partner buy-outs, trailing EBITDA—may not capture the value of firms undergoing transformation. AI maturity is difficult

to assess from the outside. Regulatory structures vary by jurisdiction and practice area. Integration of acquired firms is notoriously difficult in professional services.

1.4 The AI-Native Maturity Model

Before embarking on that journey, it is helpful for everyone in a professional services organization to evaluate where their organization stands today. Sarah needed this framework before she could build her firm—and before she left her old one.

In her final weeks at the traditional firm, Sarah created an assessment. She mapped where her old firm sat on the transformation spectrum and where she needed to go.

1.4.1 *The Five Levels*

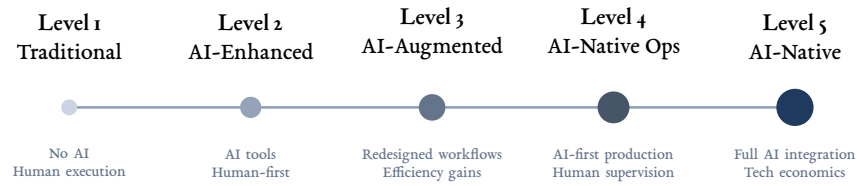

Figure 1.2: The AI-Native Maturity Model: Five levels from traditional to AI-native enterprise

The progression from Level 1 to Level 5 represents increasing integration of AI into the firm's core production system. Each level builds on the previous one, though firms can—and do—stall at any stage.

Level 1: Traditional. Sarah's old firm was here when she left. Work flowed through humans. AI was absent or experimental. Professionals billed by the hour. The staffing pyramid functioned as designed. They had a committee studying AI policy. They had done a few pilots. But nothing had changed about how work actually got done.

Many firms remain at Level 1 by choice—they serve clients who value relationships and tradition, operate in practice areas where AI has limited applicability, or lack the capital for transformation. Level 1 is not inherently wrong, but it is increasingly risky as competitors advance.

Sarah hears it constantly: "AI cannot do what I do." Sometimes this is true. Yet she would consistently observe that the person making such a claim has never tested it. They often have no idea whether the statement is actually true, either at the time or in the near future. Many have never spent more than a fleeting moment with leading AI systems. Perhaps it is simply the human condition, or something particular about those who sell expertise for a living, but knowledge workers often have a blind spot when it comes to their own disruption. They can see clearly how AI might transform other professions, other industries, other people's work—but struggle to imagine it touching their own.

Level 2: AI-Enhanced. Six months after Sarah left, her old firm moved here. They had licensed tools. Associates used Microsoft Copilot and other AI platforms for drafting and research. Some efficiency was captured. But the fundamental production system remained unchanged. Work still flowed through humans first, with AI providing support at the margins.

Most large firms today sit at Level 2. They have invested in training (to some degree), implemented guardrails, and encouraged adoption. But the economics are unchanged: hourly billing, staffing pyramid, traditional structure. Sarah recognized this was as far as her old firm would go without fundamental restructuring—restructuring the partnership would not support.

Level 3: AI-Augmented. At Level 3, the difference becomes substantive. Firms have redesigned specific workflows to incorporate AI at defined stages—a contract review process might route documents through AI classification before human review, or a research workflow might begin with AI-generated summaries that professionals verify and extend.

These firms capture real efficiency gains. They complete work faster and can handle larger volumes. But the gains remain incremental, and the core economics stay labor-based. Level 3 represents meaningful progress without fundamental transformation. Sarah saw several competitors at this level—sophisticated firms that understood AI's potential but had not yet rebuilt their operating model.

Level 4: AI-Native Operations. This was Sarah's target. Firms at this level have rebuilt their production system around AI. Work flows through AI agents first. Humans supervise, verify, and add judgment. The firm's competitive advantage comes from its AI systems and accumulated knowledge, not from the staffing pyramid.

Level 4 firms operate on different economics entirely. Margins expand as AI handles more work. Fixed-fee pricing becomes standard. The staffing ratio shifts from humans to technology. This is where the transformation from professional services firm to technology-enabled services firm becomes real.

Sarah's firm would hit Level 4 within the first year. Her contract review workflow was fully AI-native: documents flowed through specialized agents, humans reviewed outputs and handled exceptions, knowledge captured fed back into the system. Her verification process—the human supervision layer—was her competitive moat. Competitors could buy the same AI tools but the complementary operational and process layers woven together define AI-native operations.

Level 5: AI-Native Enterprise. At the apex, firms have fully realized the AI-native model. AI handles the vast majority of routine work—80 percent or more. Humans provide expert judgment on complex matters and manage client relationships. The firm's knowledge base compounds with each engagement. The economics resemble a technology company more than a traditional professional services firm.

Level 5 remains aspirational for most firms. Sarah's firm would build toward this model. Whether she achieves it, and how long it takes, is the subject we will consider in the story ahead.

Self-Assessment: Where Is Your Firm?

Level 1: Are AI tools absent from your production work?

Level 2: Do professionals use AI tools, but workflows remain unchanged?

Level 3: Have you redesigned specific workflows around AI?

Level 4: Does AI handle most work, with humans supervising?

Level 5: Does AI perform 80%+ of tasks with tech-company margins?

When Sarah assessed her old firm, they were solidly Level 1, maybe creeping toward Level 2. Sarah came to believe that the gap between Level 1 and Level 4 was too large to bridge within the existing culture of her old firm. Not without a transformative mindset and clear leadership from the top, but through her interactions with firm leadership she understood that her future lay elsewhere.

1.4.2 *Transformation Can Fail*

Sarah is on a journey. So are many other individuals and organizations. The path from one level to the next is neither linear nor guaranteed. Failure is a great teacher. And Sarah learned by watching other AI-native firms stumble and she learned through her own failures as well.

Five Risks That Destroy Transformations

Technology: AI errors create malpractice exposure.

Client rejection: Some clients pay for human attention.

Talent exodus: Your best people may leave rather than adapt.

Partnership politics: Partners near retirement may block change.

Execution: Implementations go over budget; competitors move faster.

For existing organizations, the transformation path contains many risks—the type of risks that have destroyed or impaired firms. Sarah's new venture will not be fully immune from these challenges either. How to confront and counteract those risks will go a long way toward determining the final outcome of Sarah's story.

Not every firm should transform. Not every firm that tries will succeed. Both the frameworks and the lessons from Sarah's journey can help you evaluate whether to attempt transformation and improve your odds if you do—but they cannot eliminate uncertainty.

Sarah closed her laptop at 10:10 on a Wednesday night, four months into Candor's existence. The office was empty. Elena had left at 8. Joshua had left at 6. She sat alone with a spreadsheet that was supposed to be a financial model and a coffee that had gone cold two hours ago.

She had been building the model for three hours. Revenue projections based on her pipeline. Cost assumptions based on four months of data. Margin trajectories that looked too optimistic, though she had checked the math twice. A terminal value calculation pulled from an online tutorial she did not fully trust. Half the cells were formulas. The other half were guesses dressed up as assumptions.

She pushed back from the desk and looked around the room—three desks, a shared conference table, a whiteboard covered in her handwriting, a converted Central Phoenix storefront that still smelled faintly of the sandwich shop that had occupied it before. Four months ago she had been a senior associate at a respected firm with a steady income, a clear trajectory, and a 401(k) that grew without her thinking about it. Now she had a firm name on a government form, two employees who had followed her into the unknown, a handful of clients, and a spreadsheet full of numbers she could not make tell a coherent story.

Was she crazy to have done this?

The question arrived the way it always did—not as a dramatic crisis but as a quiet voice at the end of a long day, when the work was done and the conviction

that carried her through the hours had thinned to something more fragile. She had left a career. She had convinced Elena to leave one too. She had recruited Joshua away from a firm where he was valued and comfortable. Clients were counting on her. People were counting on her.

She could not go back. That door had closed the moment she filed the paperwork—and honestly, it had closed before that, the night she sat at her kitchen table and typed ***Candor*** into a form because she saw no way forward at her old firm given the culture and leadership. So there was no version of her old life she could return to that would not feel like surrender.

She thought of that line from Robert Frost, "the best way out is always through."

So "through" it was. But she could not push through on conviction alone. She knew what Candor could be. What she could not do was explain it in the language that mattered to the people who controlled capital. Enterprise value, discount rates, terminal multiples—the alchemy by which a set of cash flow projections became a number someone would write a check against. She was a lawyer. She could argue a case or mark up a contract in her sleep. She could not look at this spreadsheet and feel confident that the story it told was the right one.

She picked up her phone and scrolled to a name she texted occasionally but had not seen in a long time. Maya Chen—an old friend. If anyone could teach her the language she needed to speak, it was Maya.

I need a crash course in the economics of what I am building. I can feel it working but I cannot model it. I am coming to New York to see you. Dinner on me. You pick the place.

The reply came before Sarah could set the phone down: *West Village. Next Thursday. And don't worry—you know I live and breathe this stuff.*

Coming Next Week

Chapter 2
The Economics of Transformation

If AI-native firms operate on different economics, what exactly does that look like? Chapter 2 builds the financial architecture of transformation: the Value Driver Tree that maps operational decisions to enterprise value, the Margin Waterfall that shows where profit expansion comes from, and the Transformation J-Curve that every firm must survive. We present the numbers behind the thesis—and the honest math that separates viable transformations from expensive failures.

Subscribe at `https://theainativefirm.com` to receive new chapters as they release.

About the Authors

Daniel Martin Katz, PhD, JD, MPP is Professor at Illinois Tech–Chicago Kent College of Law. Named by the *Financial Times* as one of the top 20 legal market shapers of the past twenty years, his research focuses on legal analytics, legal technology, and the future of the legal profession.

Michael J. Bommarito II, MSE, MA is a serial entrepreneur, researcher, and adjunct professor with twenty-five years of industry experience. He has founded and led multiple companies at the intersection of artificial intelligence and professional services.

Jillian Bommarito, CPA, CIPP/US/E is an advisor, risk and governance expert, and one of the first certified AI auditors in the world. She brings deep expertise in compliance, privacy, and the governance challenges of deploying AI in regulated industries.

Get the Full Book

`https://theainativefirm.com`