


BUILDING AI-NATIVE PROFESSIONAL SERVICES FIRMS

Strategy, Economics, and Execution



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

CHAPTER 6

Land, Expand, Compound

CHAPTER 6 · APRIL 2026

- ✓ Prologue · Four Lawyers, One Monday Morning

PART I · THE OPPORTUNITY

- ✓ Chapter 1 · The AI-Native Professional Services Firm
- ✓ Chapter 2 · The Economics of Transformation
- ✓ Chapter 3 · Uncertainty and Structure as Strategy

PART II · THE TRANSFORMATION

- ✓ Chapter 4 · Craft, Process, and Scale
- ✓ Chapter 5 · Thin Wrappers and Deep Systems
- **Chapter 6 · Land, Expand, Compound** YOU ARE HERE
- Chapter 7 · The Org Chart No One Expected
- Chapter 8 · Growing Pains: Talent and Culture
- Chapter 9 · Measure Twice, Cut Once

PART III · THE EXECUTION

- Chapter 10 · Paths—Build, Transform, Combine
- Chapter 11 · Capital and Investment
- Chapter 12 · Execution

PART IV · THE VALUE

- Chapter 13 · Value Creation and Capture
- Chapter 14 · The Future

Chapter 6: Land, Expand, Compound

*Strategy is about making choices, trade-offs;
it's about deliberately choosing to be different.*

— Michael E. Porter

The platform was functional and getting better every week. Priya had the five-layer architecture humming: retrieval, orchestration, the learning loops that improved with every engagement. David's correction pipeline was feeding data back into the system. The ContractZoo pilot had confirmed what Sarah suspected: buying someone else's wrapper was a dead end. They were building their own.

And now they were beginning to move beyond just contracts into regulatory matters and a broader range of document-intensive tasks.

But a question had been gnawing at Sarah since the early weeks, running in parallel with all the technical work, and she still did not have a clean answer: what kind of firm was she actually building? The ContractZoo experience had shown her what pure Replace looked like without depth—eighty percent of the way there, missing exactly the twenty percent clients actually paid for. The opposite extreme, treating AI as a glorified spell-checker while experienced lawyers did all the real work, left too much value on the table. What was this AI-native law firm she was calling Candor?

She could refine the existing model: make lawyers more productive with better tools and tighter processes. Or she could replace the model entirely: build a system where AI did most of the work and lawyers served as supervisors and quality gatekeepers rather than primary producers.

The choice felt binary. But it was not.

6.1 THE FALSE BINARY

David found Sarah in the office on a Saturday morning, which surprised neither of them. She had a yellow legal pad in front of her, two words written at the top—*Replace* and *Optimize*—and a look on her face that David had learned to recognize: the look of someone who was stuck and did not want to admit it.

“I thought you were taking the weekend off,” he said, setting down a coffee from Lux—his favorite CenPho coffee establishment.

“I tried. I have been reading everything I can find on AI strategy in professional services. Academic papers. Consultant frameworks. Venture capital blog posts.” She pushed the pad toward him. “Most of it is useless. The academics write about AI in the abstract. The consultants offer matrices that collapse under the weight of real decisions. The VCs want disruption narratives, not the truth about serving clients on Monday morning. What I need is simpler and harder: a theory of what this firm actually is.”

David studied the pad. “Walk me through these two options.”

“Replace means we go all in. AI does the work, lawyers supervise. The VC pitch—one lawyer does the work of a hundred. The economics are extraordinary on paper. If AI handles eighty percent, cost per matter plummets, capacity explodes. ContractZoo is a shallow version of this—layers of automation around frontier models, impressive in demos, brittle in practice. The pure version goes further.”

“And Optimize?”

“Keep lawyers at the center. Make them faster, more thorough, more productive—but they are still the ones doing the work. The PE-backed approach. Experienced lawyers, former regulators, paired with AI tools. Their

pitch is expert supervision. The value comes from who is overseeing the AI, not from the AI itself. The economics are more modest but more certain. Fifteen to twenty-five percent margin improvement without rebuilding how work gets done.”

David was quiet for a moment, running his finger along the edge of the pad. Then he said something Sarah had not expected.

“Why not just go Replace? Full commitment. You have the technology. Priya’s platform can handle it. The economics are better. And frankly, the investors want to hear that story—the one where you are building a machine, not a slightly better law firm.”

Sarah felt her jaw tighten. “Because eighty to ninety percent accuracy is not good enough when the stakes are high. A missed liability clause in a standard vendor agreement is embarrassing. A missed liability clause in a fifty-million-dollar acquisition is catastrophic. And we are not handling toy problems, David. We are handling real legal work for real companies that face real consequences when we get it wrong—and so do we.”

She leaned forward. “A traditional firm buries a missed clause in an insurance claim. An AI-native firm with an alternative business structure turns it into a regulatory case study. A large and growing percentage of our revenue comes from regulatory compliance work. If we miss something in an FDA filing or botch a state licensing requirement, the client does not just lose money—they lose their ability to operate. That is not a rounding error. That is an existential event for the client and a reputational catastrophe for us.”

“I know that. Listen, I have been through all of this type of discussion with doctors. You are letting the hard cases drive the design for the easy ones. Ninety-five percent of your contract reviews and regulatory work is actually routine. Candor OS handles them fine. You are building your entire production model around the five percent that require human judgment.”

“Because that five percent is where the malpractice claims live.”

David leaned back. “Fair. So what are you actually stuck on?”

Sarah took a breath. “Neither model fits. Replace works for volume work and fails for complex work. Optimize works for complex work and wastes ca-

capacity on volume work. I keep trying to choose between them and the honest answer is that neither one, in its pure form, describes what we are building.”

David picked up the pad and wrote a third word below the first two: *Hybrid*.

“So do not choose. Build both.”

Sarah stared at him. “That is not a strategy. That is a hedge.”

“No.” David set the pad down and drew a line separating the page into two halves. “A hedge is doing the same thing halfway. This is doing two different things completely. Replace for the volume work—contract review, compliance monitoring, standard regulatory filings. AI does seventy to eighty percent of production. A lawyer reviews the output. You compete on cost, speed, and consistency. Optimize for the complex work—M&A strategy, regulatory risk, litigation oversight. AI augments senior professionals. You compete on quality and insight. Two production systems under one roof.”

“Two production systems means two sets of workflows, two staffing models, two pricing structures. That is operationally complex.”

“I have done it before.” David’s voice carried a certainty that Sarah rarely heard from him. “In healthcare, routine diagnostics ran through a high-throughput automated system—blood work, imaging reads, standard panels. Complex cases went to specialist physicians who spent real time with the patient. Same organization, two production tracks. The trick was the triage mechanism that routed work correctly. If routine work leaked into the specialist track, you wasted expensive physician time. If complex work leaked into the automated track, you got errors. Ninety percent of the routing was algorithmic. Ten percent was human judgment on the borderline cases. The algorithm got better over time.”

Sarah saw it then—not as a compromise but as an architecture. She pulled the pad back and wrote beneath David’s line:



THE HYBRID ARCHITECTURE

Replace for high-volume, standardized work—AI handles 70–80% of production.

Optimize for complex, judgment-intensive work—AI augments senior professionals.

Smaller associate layer—AI reconstitutes the base of the staffing pyramid. Associates hired with a different set of skills.

Two production systems under one roof, integrated by a triage mechanism that routes work to the appropriate tier.

“A massive reconstitution in the associate layer,” she said, underlining it. “That is the structural break. Traditional firms need armies of junior associates to do the routine tasks. We replace some of these tasks with AI. Senior lawyers handle the judgment work. The entire base of the pyramid will be reconstituted. Ideally, further up the pyramid we need juniors to be power AI users and those who really start to understand how to identify and remediate the shortcomings of AI systems. Maybe it is not actually a pyramid but some other shape. Either way, we are running a different kind of railroad here.”

David nodded. “And the two tiers reinforce each other. Volume work generates data that improves the AI. Better AI improves the quality of the complex work. The whole thing compounds.”

Sarah’s eyes lingered on the pad. “You realize this is a bet. If the triage mechanism fails—if complex work gets routed to the volume track and we deliver garbage to a client—the whole model collapses.”

“Then we build the triage mechanism first and we build it right.” David made a note: *Build the triage function early. It is the mechanism that holds the two-tier model together.*

Sarah had her hybrid model. She had David’s operational challenge that forced her to articulate it, and his healthcare analogy that made it concrete. Now

she needed to test it against the people whose opinions mattered most: her investors (and clients).

6.2 WHERE TO PLAY, HOW TO WIN

The board call was scheduled for Thursday at noon. Sarah spent the morning reorganizing her notes three times, which she recognized as a displacement activity but could not stop doing. The yellow pad from Saturday sat on the corner of her desk, David's handwriting mixed with hers, the word *Hybrid* circled twice. She had a strategy now—or the beginning of one. What she did not have was confidence that it would survive contact with Alex Hawthorne's pattern recognition.

Alex dialed in from his home office in Los Altos. Rachel joined from Juris Dictum's conference room. David Park sat in the chair across from Sarah's desk with his laptop open to a spreadsheet he had been refining all week. Sarah noticed her hands were cold, which happened when she was nervous. She pressed them flat against the desk.

"Before we get into the quarterly update," she said, keeping her voice steady, "I want to use this call to work through something. I need your help thinking about positioning."

"Good," Alex said. "That is exactly what I was going to push you on. You have capital now. You have David building operational discipline. But I still do not have a clear answer to the most basic strategic questions: who is your customer? And what are you selling?"

The directness landed like a slap, even though Sarah had been expecting it. Alex had a habit of cutting through ceremony to the question that mattered most, and his tone carried the particular bluntness of someone who had watched enough portfolio companies drift into irrelevance while their founders avoided hard conversations.

Sarah glanced at David, who gave her a small nod. *Tell them the truth.*

"I have been operating opportunistically," she said. The word tasted like an admission. "Taking work from my network. General counsel I knew from

my old firm, a few referrals. The clients range from a twenty-person startup that needed employment agreements to a mid-size real estate developer doing a portfolio acquisition. Some of the work is good. Some of it we probably should not have taken.” She paused, then forced herself to say the part she had been rehearsing. “And we lost two prospects last month—one went back to their incumbent firm, the other decided to just use ChatGPT themselves. The margins are okay but not where they need to be. And there is no coherence to any of it.”

Saying it out loud, to her investors, with David sitting across from her, made the situation feel more real than it had when she was only thinking it.

“That is fine for the early months,” Rachel said. “You were proving the model could work at all. But you are past that now. You have capital and a shrinking runway. You need a market.”

Alex’s voice carried the particular directness Sarah had come to expect from him. “Let me be blunt. I see three options and two of them are going to be quite challenging.”

Sarah felt her stomach tighten. She had spent the weekend building a framework with David, and Alex was about to redraw the map in real time. She smiled despite herself. “Go ahead.”

“Option one: enterprise. Fortune 500 in-house legal departments. The Cisco and Walmart and JP Morgan Chase type-buyers. Big budgets, sophisticated procurement, large deal sizes.” He paused. “The problem is that enterprise sales cycles are twelve months minimum. Often longer. You will burn through half your runway before you close your first real enterprise deal. You do not have the brand, the track record, or the sales infrastructure to compete for that work yet.”

He let that sink in. “The well-funded AI-native firms can do it because they raised ten to hundreds of millions of dollars and they have former Am Law 100 partners opening doors. You raised three million. Feel free to be opportunistic if a clear opportunity presents itself, but in general, enterprise is a later-stage play.”

“Option two,” Rachel added. “Small business and consumer. The LegalZoom market. High volume, low price, massive addressable market.”

“The problem there is unit economics,” Alex said. “Customer acquisition cost for small business legal services is brutal. You are spending two hundred to three hundred dollars in marketing to acquire a client who pays you five hundred dollars for a one-time engagement. The math only works at enormous scale, and you do not have the marketing budget or the brand recognition to reach that scale. That is a venture-backed platform play, not a professional services firm play. LegalZoom spent years and hundreds of millions building that acquisition engine. They literally have NBA commercials with Giannis. You cannot replicate it with three million dollars.”

“Which leaves option three,” Sarah said.

THE MID-MARKET SWEET SPOT

Companies with 50–500 employees and \$500K–\$5M+ in annual legal spend. Sophisticated enough to evaluate an AI-native value proposition. Frustrated enough with traditional delivery to try something new. Sales cycles of 4–8 weeks, not 6–12 months. Deal sizes of \$40K–\$150K per engagement. Genuinely underserved: too small for elite firms, too complex for self-service platforms.

“Mid-market.” Alex’s tone shifted from diagnostic to prescriptive. “Companies with at least some degree of scale. The ones with a general counsel—or maybe a VP of Legal who is really a business executive wearing a legal hat—who is spending half a million to five million dollars a year on outside counsel and hating every invoice. These buyers are sophisticated enough to evaluate your value proposition but not so sophisticated that they need a twelve-month procurement process to approve a new vendor. They are price-sensitive enough that your cost advantage matters but not so price-sensitive that they will not pay for quality. And there are tens of thousands of them.”

Sarah’s hands were warm again. She had been thinking along similar lines, but hearing Alex articulate it—with the conviction of someone who had helped

build companies in this exact market segment—turned her private intuition into something that felt like shared conviction.

“The mid-market GC,” Alex continued, “is the most underserved buyer in legal services. Think about it. Big firms do not want them because the deal sizes are too small to justify partner attention. A mid-market company paying a million or two a year in total external legal fees barely registers at an Am Law 100 firm. Small firms cannot serve them properly: they lack the capacity, the technology, and often the breadth of expertise. ALSPs have tried to capture this market, but they carry the stigma Rachel mentioned on our hike—the buyer is a lawyer, and lawyers are skeptical of non-law-firm providers. And legal tech platforms sell tools, not outcomes. The mid-market GC does not want another piece of software to manage.”

He was building momentum now. “What mid-market GCs really want is a law firm that operates like a technology company—fast, transparent, predictable on price, and capable of handling a range of corporate legal needs without requiring a stable of different outside counsel for each issue. But make no mistake, it is not really the tech they are buying. It is the outcomes, timeliness, and cost effectiveness they are purchasing. The tech is an important input, but your customer is only partially interested in your internal division of labor and machines.”

“That is exactly what Candor is building,” Sarah said.

“Then say so. Explicitly. Relentlessly. Build your entire go-to-market around that buyer. Every piece of marketing, every sales conversation, every service design decision should start with the question: does this serve the mid-market GC who is spending half a million to five million on outside counsel?”

Sarah spent the next week testing phrases on David, on Priya, on the whiteboard in the conference room. Most of them were too long, too clever, or too generic—the kind of language that sounded good in a pitch deck but collapsed under the weight of a real conversation. The version that stuck was the one David scribbled on a Post-it during a late evening working session and slapped on her monitor without comment: *AI Powered. Lawyer Delivered. Trust Assured.* Three phrases. No jargon. The first told you how the work got done. The second told you who stood behind it. The third told you why it mattered. Sarah left the Post-it on her monitor for months.



Figure 6.1: Alternative Logo Distilling Candor’s Value Proposition for the Mid-Market General Counsel

David, who had been listening quietly, leaned forward. “Can I add something from the operations side?”

“Please,” Alex said.

“The mid-market fits our production model better than either extreme. Enterprise clients want bespoke service—custom workflows, dedicated teams, white-glove treatment. That is expensive to deliver and hard to standardize. Every enterprise engagement becomes a one-off project. Small business clients want self-service—they do not want to talk to a human at all if they can avoid it. That requires a consumer-facing platform we have not built and that requires a completely different set of skills to operate. Mid-market clients want something in between: professional service quality with technology efficiency. They want a human they can call when things get complicated, but they also want their more straightforward or commodity work done in hours to days instead of weeks to months. That is the sweet spot for our hybrid model—human enough to feel like a law firm, automated enough to deliver like a technology company.”

Rachel jumped in. “David, what does the data say about the economics? If you target mid-market, what does the average engagement look like?”

David pulled up his spreadsheet. “Revenue run rate is about \$1.6 million from six ongoing clients. Three of those are expanding scope—Linda Torres just sent us a second batch of contracts, and two others have moved from one-off engagements to recurring work. The other three were one-offs that probably will not repeat. Blended gross margins are forty-five percent, EBITDA around twenty-eight percent. Lower than I want—we are still eating rework costs as the platform matures. But the trend is clear: every engagement reduces the

verification burden on the next one. Time is on our side and greater margins are coming.”

“What about the revenue concentration risk?” Rachel asked. Her voice had shifted from strategic to fiduciary—the tone of an investor who needed to understand how fragile the foundation was. “Six clients is not a market. It is a contact list. And your largest client is what—forty percent of revenue?”

“Closer to thirty-five,” David said. “But the point stands.”

Sarah felt the familiar tightness return. She could have softened this, positioned it as a growth opportunity rather than a vulnerability. But that was not what the moment required.

“Agreed,” she said. “That is why we need a deliberate acquisition strategy. And honestly, winning new clients has been harder than I expected.” She heard the wobble in her own voice and pushed through it. “We have had prospects go dark on us, prospects who loved the pitch but could not get budget approval, prospects who went back to their incumbent firm because switching felt too risky. One prospect told me to my face that she loved what we were building but could not justify the career risk of hiring a startup law firm.” She let that sit for a moment. “Alex’s point about the mid-market gives us a target. David’s economics give us a value proposition. But we need to figure out how to find and convert twenty more folks like Linda Torres in the next twelve months.”



After the operations discussion, the call turned to a question that Sarah had been wrestling with privately: how to position the firm’s two tiers of service in a way that was coherent to the market.

“You are describing two different businesses,” Rachel said. “One is a high-volume, AI-driven production operation that handles routine legal work at scale. The other is a boutique advisory practice where experienced lawyers provide strategic counsel. How do you sell both without confusing the market?”

Sarah had sketched the answer on her yellow pad three days earlier. She pulled it up now.

“Think of it as one firm with two service tiers. Tier one is what I am calling our Standard offering—contract review, regulatory compliance monitoring, due diligence document analysis, standard corporate filings. This is AI-native volume work. AI does seventy to eighty percent of the production. A senior lawyer reviews the output, catches errors, handles edge cases. We price this on fixed fees, per-unit or per-engagement. We compete on cost, speed, and consistency.”

“And tier two?” Alex asked.

“Tier two is our Advisory offering. Complex matters that require genuine professional judgment. M&A strategy, regulatory risk assessment, litigation strategy, corporate governance. This is AI-augmented advisory work. AI does the research, the document analysis, the preliminary assessment. But a senior lawyer drives the analysis, develops the strategy, and advises the client. We price this on value—enhanced fixed fees or scoped engagements, not hourly billing, but at premium rates that reflect the expertise involved.”

“So tier one is your Replace model and tier two is your Optimize model,” Rachel said.

“Exactly. And the key insight is that they reinforce each other. Tier one generates volume, which builds our data advantage and funds the technology investment. Tier two generates margin, which funds senior talent and builds client relationships. A mid-market GC who uses us for contract review at tier one starts to trust us. When they have a complex issue—an acquisition, a regulatory investigation, a governance question—they come to us for tier two advisory instead of going to a big firm they barely know. We capture more of their legal spend over time. The land-and-expand playbook that every SaaS company runs—except we are running it with legal services.”

Alex was quiet for a moment. Sarah could hear him thinking, which she had learned was different from him being silent. When Alex was silent, he was processing. When he was thinking, she could almost hear the gears turning, testing her logic against his pattern recognition from decades of deals.

“That is a good story,” he said finally. “But it creates a positioning challenge. In the market’s mind, are you the cheap option or the expert option? Because those are usually different firms. The general counsel who hires you for cut-rate contract review may not trust you for bet-the-company advisory. And the GC

who values your senior expertise may wonder why you are also running a volume shop. Perception matters. Brands that try to be both premium and value tend to be neither.”

Sarah had anticipated this objection. “The analogy I keep coming back to is the Big Four accounting firms. Deloitte runs a massive audit practice—high-volume, process-driven, increasingly technology-enabled. They also run a premium consulting practice—bespoke, relationship-driven, expertise-intensive. Nobody thinks it is strange that the same firm does both. The audit work creates client relationships and institutional knowledge. The consulting work generates premium margin and deepens trust. They reinforce each other.”

“The Big Four had decades to build that dual positioning,” Rachel said. “And they started with audit as a regulatory mandate that guaranteed deal flow. That is a built-in client acquisition engine you do not have. Deloitte does not need to convince anyone to buy an audit—the law requires it. You need to convince every single client to take a chance on a startup. The analogy is aspirational, not structural.”

Sarah absorbed the hit. Rachel was right that the structural advantage was different. But the principle held.

“True. But I do not need Deloitte’s scale. I need twenty to thirty mid-market clients who see the value of both tiers. And the beauty of the mid-market is that these are the clients who *need* both. We just need more than one entry point for a client to connect to us.” She let that settle. “A company with two hundred employees and a million-dollar legal budget cannot afford one firm for routine work and a different firm for complex matters. They are not going to run an RFP every time they need a different type of legal service. They want one relationship. One firm that knows their business, understands their risk tolerance, and can flex between routine execution and strategic advice. We offer that, at a price point that makes sense for their budget.”

⚠ THE BRAND CONFUSION RISK

Firms that position as both premium and value tend to be neither. The general counsel who hires you for cut-rate contract review may not trust you for bet-the-company advisory. Managing dual positioning requires explicit tier separation: distinct pricing, distinct deliverables, distinct SLAs. The tiers must reinforce each other through a flywheel—volume work builds data advantage, advisory work builds margin and trust—rather than competing for the same brand perception.

“We manage it through the tier structure itself. Standard tier has its own pricing, its own deliverables, its own SLAs. Advisory tier has its own. The client understands the difference because the difference is explicit—it is in the engagement letter, in the pricing, in the staffing. We are not pretending that contract review and M&A advisory are the same thing. We are saying they are different capabilities from the same firm, integrated by a shared understanding of the client’s business. That is a feature, not a bug.”

Alex paused again, then said something Sarah did not expect. “I think you are right. And I think the firms that will struggle are the ones that try to occupy only one quadrant. Pure volume players will commoditize. Pure advisory players will be undercut. The winners will be the ones who can do both—and who build the operational discipline to keep the two tiers from contaminating each other. David, that is your job. Do not let the volume work get sloppy and do not let the advisory work get bureaucratic. Now this all should be subject to revision as things progress and if our shared view turns out to be wrong then you can simply calibrate accordingly.”

“Understood,” David said, making a note.

Priya, who had been listening on mute while debugging a retrieval issue, unmuted. “One thing from the engineering side. The platform can handle both tiers—the architecture is tier-agnostic. But the Standard workflows need to be prioritized because that is where the data flywheel spins fastest. Every contract

we process through the Standard tier makes the Advisory tier smarter too. The learning compounds across both. So sequence matters.”

“Noted,” Sarah said. “Standard first, Advisory second. Build the data engine on volume, then apply it to complexity.”



The following Tuesday, Sarah and David spread his spreadsheets across the conference table and tried to solve pricing. David had built three models overnight, each with different assumptions about AI compute costs, reviewer hours per contract, and overhead allocation. Sarah had brought her own data: the invoices she had collected from prospects and existing clients, showing what mid-market companies were actually paying their current firms.

“The Standard tier is the one that matters first,” David said. “If we price it too high, we lose the cost advantage that gets us in the door. If we price it too low, the margins collapse and we cannot fund the platform.”

He pulled up the first spreadsheet. “Our direct cost per contract on a standard review—AI compute, reviewer time at current efficiency, quality sampling, delivery formatting—is running about sixty-eight dollars. That number will come down as the platform improves, but today, sixty-eight is the floor.”

“And traditional firms are charging what?”

“Brennan quoted Linda three hundred per contract. The numbers I have seen from prospects range from two hundred to three-fifty, depending on the firm and the complexity mix. The ALSPs come in lower—around one-fifty to two hundred—but they are staffing with contract attorneys in cheaper markets, not actually automating.”

Sarah studied the numbers. “If we price at one-fifty per contract, our gross margin is about fifty-five percent at current efficiency. As the system improves and direct costs drop toward fifty dollars, margin climbs toward sixty-seven percent. That is healthy.”

“And one-fifty is low enough to be compelling against traditional firms but high enough that it does not look like we are giving it away. There is a credibility problem if you price too low. A general counsel who sees fifty dollars per

contract is going to wonder what corners you are cutting. One-fifty says we are meaningfully cheaper, not suspiciously cheap.”

Sarah nodded. “What about volume discounts? If someone sends us five hundred contracts, do we go lower?”

“Not on price per unit. We go faster. The cost advantage of volume is in our utilization—a five-hundred-contract batch spreads the fixed setup cost across more units, which improves our margin without cutting the price. If we start discounting per unit, we train clients to negotiate, and we erode the whole model. We compete on speed and quality at that price point, not on further discounts.”

They turned to the Advisory tier. This was harder. Advisory engagements varied in scope, complexity, and the amount of senior lawyer time required.

“The challenge,” David said, “is that advisory pricing has to signal expertise without scaring a mid-market budget. A general counsel spending one-point-two million a year is not going to approve a two-hundred-thousand-dollar advisory engagement. But if we price too low, we undervalue the work and we attract the wrong expectations.”

Sarah thought about what a buyer like Linda Torres could actually absorb. “Scoped engagements. Not hourly, not open-ended. A regulatory compliance assessment for a defined scope—evaluating a company’s vendor agreements against a specific regulatory framework—priced at thirty-five to fifty thousand dollars depending on the number of agreements and complexity. M&A due diligence support for a defined transaction, sixty to ninety thousand. Each engagement has a defined deliverable, a defined timeline, and a defined price. The client knows what they are buying before they sign.”

David ran the margins. “At those price points, with a senior lawyer spending forty to sixty hours and the AI handling research and preliminary analysis, we are looking at sixty to sixty-five percent gross margins. Better than Standard, because the human time is higher-value and the AI contribution on the research side keeps total hours manageable.”

“Good,” Sarah said. “One-fifty per contract for Standard. Scoped engagements starting at thirty-five thousand for Advisory. Both fixed-price, both with defined deliverables, both with turnaround commitments.” She wrote the numbers on the whiteboard. “Now we need someone to actually say yes.”

The board call continued for another forty minutes, shifting from strategy to tactics. Alex pushed Sarah on competitive positioning—who exactly would she be taking share from, and why would they not respond?

“Walk me through the field,” he said. “Who loses when you win?”

Sarah had thought about this. “Traditional mid-market firms have the relationships but their cost structures are thirty to fifty percent higher than ours. They are built on associate labor and hourly billing—both of which we undercut structurally. ALSPs promised efficiency for years but most of them are built on labor arbitrage, not genuine technology leverage. They offshore the work to cheaper lawyers instead of actually automating it. Quality of work is highly varied and is sometimes downright terrible. Clients have figured that out, and the skepticism creates an opening for a firm that can demonstrate real AI capability rather than just marketing it.”

“And the Big Four?” Rachel asked.

“Long-term threat, not immediate. KPMG has an Arizona ABS but they are focused on enterprise compliance, not mid-market legal services. They will matter eventually. Not yet.”

Alex raised one more point. “There is a professional responsibility angle here that most people are not talking about yet. The ethical question used to be whether it was acceptable to use AI in legal practice. That question is inverting. Pretty soon the question will be whether it is professionally defensible *not* to use it. When an AI system catches a cascading cross-reference that a tired associate missed at hour six—or completes an analysis in four hours that would have taken forty—the general counsel who chose the traditional approach is going to have a hard time explaining that choice. The firms that adopt AI with proper safeguards are not taking a risk. They are managing a risk that is becoming harder to avoid. That is a tailwind for you.”

Rachel probed the unit economics of each service tier, asking David to model scenarios where volume fell short of projections or where advisory work proved harder to sell than expected.

When the call ended, Sarah sat in her office, the Phoenix sun slanting through the blinds, and thought about what had just happened. She had walked into the call with an intuition and walked out with a strategy. Not because Alex or Rachel had told her what to do—they had pushed and questioned and challenged, but the choices were hers. The mid-market focus. The two-tier service architecture. The hybrid operating model. The recalibration of the associate layer. Each decision narrowed her options and clarified her path.

She picked up the yellow pad and drew a simple two-by-two matrix. The horizontal axis was AI intensity: how much of the work does AI perform? The vertical axis was service complexity: how much human judgment does the work require?

Four quadrants emerged. High AI, low complexity: *AI-Native Volume*, her Standard tier. High AI, high complexity: *AI-Augmented Advisory*, her Advisory tier. Low AI, high complexity: *Traditional Premium*, the elite boutiques and Am Law firms. Low AI, low complexity: *Commodity*, the danger zone where traditional mid-market firms sat, delivering standardized work through manual processes at prices that AI-native competitors would soon undercut.

	AI Intensity	
	High	Low
Service Complexity: High	AI-AUGMENTED ADVISORY Candor Tier 2	TRADITIONAL PREMIUM Elite boutiques
Service Complexity: Low	AI-NATIVE VOLUME Candor Tier 1	COMMODITY Vulnerable to disruption

Table 6.1: AI-Era Positioning Matrix

Candor would occupy two quadrants simultaneously—and that was the strategy, not a contradiction. The two positions reinforced each other through a flywheel: volume work generated data that improved the AI, which improved

advisory quality; advisory work generated margin that funded technology investment, which improved volume efficiency. Competitors who started later would face a capability gap that widened with every engagement.

She thought about what made that position defensible. Four things, really—and each one was grounded in something Candor had already built, not something it planned to build someday.

Data advantage first. The correction pipeline David and Priya had constructed was not a theoretical asset—it held thousands of structured annotations from every engagement the firm had completed. Every time a reviewer caught a misclassified indemnification clause or adjusted a risk rating the AI had scored incorrectly, that correction entered the knowledge base with structured tags: clause type, contract category, error pattern, client context. The system that had seen a thousand variations of indemnification clauses recognized the unusual ones instantly.

A competitor launching today with the same foundation models would produce competent output on day one. But without the accumulated correction data, without knowing that procurement agreements in the healthcare sector routinely contained non-standard data handling provisions that interacted with BAA requirements in ways the base models missed, their output would remain generic. Data moats could not be built retroactively. You could not go back and capture the corrections you never logged, the edge cases you never tagged, the client preferences you never structured. The advantage accrued only to firms that designed for data capture from the beginning, and it compounded with every engagement completed.

Workflow advantage second. David's process engineering produced consistent quality regardless of which lawyer was reviewing the output. The review protocol, the confidence thresholds, the sampling methodology, the escalation criteria—all of it was documented, calibrated, and continuously refined. A new reviewer could reach baseline competence within days because the system encoded the judgment that would otherwise take months to develop through apprenticeship.

Talent advantage third. In a world where AI handled the reading, extracting, and classifying, the quality of human oversight became the primary differentia-

tor. Candor's senior reviewers were not checking boxes—they were applying judgment shaped by years of practice to the specific cases where AI reached its limits. The firm did not need an army of associates. It needed a smaller number of experienced lawyers whose expertise was amplified by the platform rather than diluted across routine tasks.

And brand—the hardest to build and the easiest to destroy. Clients needed to trust their providers more than ever, because they could not easily evaluate AI quality themselves. A track record of successful engagements and transparent reporting, including the willingness to document limitations as Sarah had done with Linda Torres, served as a proxy for quality assurance. It told the client that someone competent and honest was supervising the machine.

She turned to David, who was still sitting across from her, making notes.

“What do you think?” she asked.

David looked up. “I think the strategy is sound. The mid-market target makes sense for our production model. The two tiers map onto different workflow architectures, which means I can tune each one independently.” He paused. “And now I know what to tell Priya. She has been asking me which workflows to prioritize in the platform build. The answer is the Standard tier first—contract review, compliance monitoring, the volume work. That is where the mid-market clients will enter. Once we prove the economics there, we layer in the Advisory workflows.”

“Agreed. The five-layer architecture she designed can handle both tiers. We just need to sequence the engineering work to match the go-to-market.”

“I will sit down with her Monday and map the sprint priorities against the mid-market requirements. She mentioned yesterday that the workflow tooling is maturing fast—she used Claude Code to build the initial healthcare compliance pipeline in a fraction of the time the contract review pipeline took. Alex's push on using AI to build AI is starting to pay off. The data capture pipeline she and I have been building? It is exactly what these clients need—every engagement feeding back into the system, compounding the advantage. Now we have a clear target to aim it at.”

Sarah nodded. She had the strategy, the target market, and a team that understood both. What she did not have yet was a client who had actually said yes.

6.3 A BUYER IN WAITING

Two weeks after the board call, Sarah was learning the difference between having a strategy and having clients who cared about it.

The first prospect, a VP of Legal at a logistics company in Tucson, listened politely for thirty minutes, asked several sharp questions about the technology, and then said he was “not quite ready to make a change but would love to stay in touch.” Sarah recognized the language from her years at the old firm—it was a no wrapped in courtesy. She sent a follow-up email that went unanswered.

The second prospect never showed for their video call. Sarah sat in her office for twelve minutes, refreshed the meeting link twice, then sent a brief “hope everything is all right” note. The reply came three days later: *“Apologies—got pulled into something. Let’s reschedule.”* They never did.

The third was the one that stung. Mark Hensley ran legal for a chain of urgent care clinics in the Southwest—exactly the profile Alex had described, exactly the buyer the strategy was designed for. Sarah and David had spent two full days preparing for the pitch. David built a custom invoice comparison: Hensley’s last twelve months of outside counsel bills, line by line, mapped against what Candor would have charged for the same work. The savings were stark: \$340,000 in spend that could have been \$195,000.

Sarah had Priya mock up a workflow dashboard showing what Hensley’s contract intake would look like inside Candor OS—documents flowing through the pipeline, confidence scores updating in real time, a compliance tracker tailored to healthcare regulatory requirements. It was the most thorough pitch package they had ever assembled.

Hensley was enthusiastic in the first meeting, asked David to walk him through the workflow in detail, even shared his current outside counsel invoices so they could sharpen the comparison. Then, two days before the engagement letter was supposed to go out, he called Sarah directly.

“I have been thinking about this all weekend,” he said. “And I just cannot get there. Not yet. You are asking me to bet my legal function on a firm called Candor that has been operating for less than a year. If something goes wrong—if a contract review misses something material, if there is a data breach, if the AI makes an error that ends up in front of a regulator—I am the one who has to explain to my CEO why I hired a startup instead of a real firm. Call me in six months, when you have more of a track record. I want to do this. I am just not brave enough yet.”

Sarah thanked him, hung up, and sat very still for a long time.

David found her in the conference room twenty minutes later, staring at the invoice comparison spreadsheet still open on her laptop.

“We did everything right on that one,” she said. Her voice was flat. “The prep work, the custom demo, the pricing. What did we miss?”

David sat down across from her and folded his hands on the table. “We did not miss anything. We are not doing anything wrong. We are doing early.”

“What does that mean?”

“It means the pitch is sound. The pricing is sound. The product is real. But you are asking someone to be first. And nobody wants to be first, Sarah. They want to be second. They want to be the person who hires you after someone else proved it works. After someone else took the career risk and came out fine. Mark Hensley is not saying no to Candor. He is saying he does not want to be your proof point. He wants to use someone else’s proof point to justify the decision.”

Sarah pushed back. “So what—we just wait? We sit here until someone accidentally stumbles into being first and then we use them to convince everyone else?”

“No. We find the buyer who does not think of it as going first. Someone whose pain is bad enough that the risk of staying with what they have is worse than the risk of trying something new. That buyer is out there. We just have not found her yet.”

The strategy was right. The positioning was right. The economics were right. And none of it mattered if the buyer could not get past the risk of being first.

A week after the Hensley rejection, Sarah was on the phone with Maya Chen. It was one of their irregular catch-up calls—Maya checking on her investment referral, Sarah downloading the emotional weight of the past month. Maya mentioned, almost in passing, that she had heard something interesting from a client.

“Do you know an ALSP called Lexicon Solutions?” Maya asked.

“I have heard the name. Contract review, compliance support, that kind of thing. They have been around for a few years.”

“One of my clients—a mid-market manufacturer, exactly your target profile—hired them for a vendor contract review. Two hundred agreements. Lexicon pitched it as AI-powered analysis, the whole thing. Fast turnaround, lower cost. The general counsel was excited.”

“And?”

“The output was mediocre. The AI component turned out to be a first-pass screen that flagged obvious deviations—missing signature blocks, expired terms, that sort of thing. Everything substantive was handled by contract attorneys working out of a shared services center in Hyderabad. The analysis memos read like they were written by people who had never seen the client’s business before, because they had not. One memo flagged a standard limitation of liability cap as ‘potentially problematic’ without any analysis of whether the cap was appropriate for the contract value and risk profile. Another missed a non-compete provision buried in a services addendum entirely.” Maya paused. “The general counsel told me she felt like she had paid for technology and received labor arbitrage in a slightly nicer package.”

Sarah felt a complicated mix of validation and frustration. Validation because the experience confirmed exactly what she had been telling prospects: most firms marketing AI-enhanced services were wrapping old delivery models in new language. Frustration because every bad experience with a competitor made the next sales conversation harder. Lexicon’s failure did not make prospects more willing to try Candor. It made them more skeptical of anyone claiming to use AI at all.

“There is something else,” Maya continued. “I also heard that Henson and Locke—the mid-size firm in Denver—is telling clients they offer ‘AI-enhanced’ contract review now. You know what their AI enhancement turned out to be? They gave their associates access to ChatGPT and told them to use it for first drafts. No quality system. No confidence scoring. No structured workflow. Just associates pasting clauses into a chat window and copying the output into memos. One of their associates told a friend of mine that he has no idea whether the AI output is right, so he just rewrites most of it anyway. The firm charges the same hourly rates and calls it innovation.”

Sarah sat with that for a moment. The competitive field was not empty—it was crowded with firms that had adopted the language of AI without building any of the infrastructure that made AI actually reliable. That created both a problem and an opportunity. The problem was that buyers could not easily distinguish between genuine AI-native capability and marketing. The opportunity was that the firms doing it badly were creating dissatisfied clients who would eventually look for something better.

“The bar is low,” Sarah said.

“The bar is underground,” Maya replied. “Which means when someone actually clears it, the difference will be obvious. You just need to get enough people to let you prove it.”



Then Alex called with a Denver introduction, and Sarah booked a flight before she could talk herself into waiting. The prospective client was a general counsel named Allison McLindon, who ran legal for a regional healthcare company. Four hundred employees, twelve clinics across Colorado and Utah. Annual external legal budget: two million dollars, split among four outside firms.

“I hate it,” Allison said, stirring her coffee. They were in a hotel lobby near the Denver airport, the kind of place where business travelers met because it was convenient for no one and therefore equidistant for everyone. The chairs were too close together and the ambient noise was relentless—a constant stream of rolling suitcases, hotel checkouts and the hiss of a poorly calibrated espresso

machine. “Four firms, four billing systems, four sets of engagement letters, four year-end audits of outside counsel spend. None of them talk to each other. I spend twenty percent of my time managing outside counsel instead of doing actual legal work for the company.”

“What would you change if you could?” Sarah asked.

“Fewer firms. A deeper relationship. Particular focus on someone who can handle the routine stuff—contracts, compliance, standard corporate filings—while bringing in heavy expertise when I need it. An acquisition, a regulatory investigation, a complicated employment dispute. I do not need a hundred-lawyer firm. I need a smart, responsive firm that uses technology the way I use it in every other part of my business. My clinics run on electronic health records that talk to our billing system that talks to our scheduling platform. But my legal function operates like it is 1995. Email, Word documents, invoices that show up three months after the work is done.”

Sarah recognized the buyer Alex had described. The mid-market GC who was spending over a million dollars on outside counsel and getting mediocre service. Not because the lawyers were bad—they were perfectly competent. But because the delivery model was designed for a different era. Four separate firms, each billing hourly, each treating Allison’s matters as bespoke projects, each staffing with associates who were learning on her dime.

“Let me tell you what Candor does,” Sarah said, and walked Allison through the two-tier model. Standard services: contract review, compliance monitoring, corporate filings, all delivered through AI-native workflows at fixed prices with guaranteed turnaround times. Advisory services: M&A support, regulatory strategy, litigation oversight, delivered by experienced lawyers augmented by AI on a scoped engagement basis with clear deliverables.

“What does that look like in practice?” Allison asked. “Say I have three hundred vendor contracts that need to be reviewed for a new compliance requirement. HHS just updated the requirements for business associate agreements, and I need to know which of my vendor contracts are compliant and which need to be renegotiated.”

“Under the Standard tier, we ingest all three hundred contracts into our analysis system. AI identifies the relevant clauses, flags potential compliance gaps

against the new HHS requirements, and generates a structured report for each contract: compliant, non-compliant, or requires further review. A senior lawyer who specializes in healthcare regulatory work reviews the AI output, validates the analysis, and handles any edge cases or ambiguities. You get a complete compliance assessment—three hundred contracts, each with a detailed report and a recommended action—in five to seven business days. Fixed price. No surprises on the invoice.”

“How much?”

“For three hundred standard vendor contracts against a defined compliance standard, roughly forty-two thousand dollars. Call it a hundred and forty per contract.”

ALLISON McLINDON: THE TEMPLATE CLIENT

Regional healthcare company. 400 employees, 12 clinics across Colorado and Utah. Annual external legal budget: \$2M, split among four outside firms. Pain points: four billing systems, four engagement letters, 20% of her time managing outside counsel. Consolidation opportunity: \$800K in annual spend that could move under one roof. First engagement: 300 vendor contracts, \$42K fixed fee, 7-day turnaround. If this went well, it would be the proof point Sarah desperately needed. If it didn't, it would confirm every skeptic who said a startup law firm couldn't deliver at scale.

Allison's eyebrows rose. “My current firm quoted me over one hundred thousand for the same project. They estimated three associates working for three weeks.”

“That is the production model difference. Their four associates read documents page by page. Our AI processes all three hundred contracts in parallel and a senior lawyer supervises the output. We deliver the same analytical quality—arguably more consistent quality, because AI does not get tired at hour six or distracted by a text message at hour eight—at a fraction of the cost and in a fraction of the time.”

“And if one of those contracts has a complex indemnification issue that the AI flags? Something non-standard that requires real analysis?”

“That is where the model shines. The AI flags it. The senior lawyer reviews it and determines it needs deeper analysis. We shift that specific issue to our Advisory tier. An experienced lawyer—someone with fifteen or twenty years of practice, who has actually negotiated these indemnification structures—analyzes the provision, assesses the risk exposure, and advises you on how to proceed. Do you renegotiate? Do you accept the risk? Do you need to restructure the relationship? That analysis is priced separately, on a scoped basis, at a rate that reflects the expertise involved. You are not paying for a junior associate to research the issue from scratch. You are paying for someone who already knows the answer and is using AI to verify and document it.”

Allison was quiet for a moment. “So the routine work is cheap and fast, and the complex work gets real expertise instead of a third-year associate Googling indemnification clauses at two in the morning.”

Sarah laughed. “That is a more honest description than anything in our marketing materials.”

“Can I see the AI output? I mean, can you show me what the analysis actually looks like?”

Sarah opened her laptop and pulled up a sample report from a recent engagement, anonymized but structurally identical to what Allison’s contracts would produce. She scrolled to a specific entry.

“This is a food services vendor agreement. The AI flagged it amber—the indemnification cap is uncapped, which is unusual for this contract type, and the auto-renewal clause limits termination to a thirty-day window that most companies miss. The confidence score on the indemnification finding is ninety-four percent, meaning the AI is highly certain. The auto-renewal flag is eighty-seven percent, lower because the termination window interacts with a separate exhibit that modifies the renewal terms. That interaction is exactly the kind of thing a junior associate skimming at hour six might miss. The AI catches it because it reads every clause against every other clause, every time, without fatigue.”

Allison leaned forward, studying the screen. She scrolled through several more entries on her own, pausing at the summary dashboard—how many contracts were compliant, how many needed attention, where the highest risk concentrations sat.

“This is better than what I get from my current firm,” she said. “More structured. More consistent. I can actually compare across contracts instead of reading fifty different associate-authored memos written in fifty different styles. And this dashboard at the top—I could take this straight to my CEO and explain the compliance picture without having to translate from lawyer to English.”

“That is the point. When AI produces the first draft, the output is consistent by design. The structure, the terminology, the level of detail—all standardized. Human review adds the judgment layer on top. The combination is more thorough and more usable than either alone.”

Allison closed the laptop but did not say yes immediately. “I need to think about this over the weekend. And I want to call a reference—someone who has actually used your firm for a real engagement, not a demo.”

“Of course.” Sarah had anticipated this. “I will connect you with Linda Torres. She ran three hundred and fifty contracts through our system two months ago. She will tell you what worked and what did not—including the parts we got wrong on the first pass.”

Sarah hesitated, then decided that a firm called Candor should act like one. “I should also be transparent: healthcare regulatory compliance is a newer domain for us. We have configured the pipeline for HHS-specific requirements, and Priya, our CTO, has built the cross-reference logic for business associate agreements. But this will be our first large-scale healthcare engagement. We will staff it with a senior lawyer who has healthcare regulatory experience, and we will over-invest in quality review on this one. If the output does not meet your standards, you do not pay.”

Allison studied her for a moment. “I appreciate the honesty. That is actually more reassuring than if you had told me you had done a thousand of these. It is a brand new requirement anyway.” She paused. “Send me the engagement letter. Let me talk to your reference over the weekend. If Linda confirms what you have shown me, we will start with the vendor contract review and see how it

goes. If that works the way you described, I have about eight hundred thousand dollars of annual legal spend that I would love to move under one roof. But one step at a time.”



The weekend after Denver was one of the longest of Sarah’s life. Data was not the only flywheel that could power Candor. Word of mouth referrals and success stories that clients were willing to share could allow Sarah’s AI-native law firm to achieve accelerated growth. She had given Allison Linda Torres’s number on Friday afternoon, and now there was nothing to do but wait.

Sarah tried to work on Saturday morning—reviewing Priya’s sprint plan, editing the engagement letter template David had refined—but her concentration kept fracturing. She would find herself staring at her phone, willing it to produce a text from Allison, then feeling foolish for staring at her phone, then picking it up again two minutes later.

She thought about what Linda would say. The Torres engagement had been real. Not a curated demo, not a sanitized case study—a real engagement with real problems. The first batch had come back with an eighteen percent error rate that still made Sarah’s stomach clench when she thought about it. The delivery had been two days late. Linda’s callback had been measured, not warm.

She had found contracts with mismatched risk ratings, inconsistent formatting, and what looked like copy-paste analysis on three high-risk agreements. Those were facts. They were in the record. And Sarah had given Allison permission to hear all of it.

David had told her this was the right move. “If Linda sugarcoats it, Allison will not trust the reference. If Linda is honest about the problems and also honest about the recovery, that is more credible than a perfect story. Nobody believes perfect.”

Sarah knew he was right. She also knew that the outcome was entirely out of her hands. Linda Torres would say whatever Linda Torres decided to say, and Sarah could not script it, coach it, or control it.

Sunday passed. No word. Sarah went for a run along the canal paths near Tempe Town Lake, pushing harder than usual, trying to burn off the nervous energy. She got home, showered, checked her phone. Nothing.

Monday morning arrived. Sarah was at the office by six-thirty, pretending to review a compliance assessment for an existing client. At nine-fourteen, her phone buzzed.

A text from Allison McLindon: *“Talked to Linda Torres yesterday. She did not hold back—told me about the error rate on the first batch, the late delivery, the formatting problems. She also told me you sent her a cover memo documenting every limitation before she found them herself, which she said no firm has ever done. She said the second batch was meaningfully better—faster turnaround, cleaner analysis, the risk ratings were consistent. She is sending you more work. Her exact words were: ‘They are not perfect, but they are honest about it, and they are getting better fast. That is more than I can say for any firm I have used in twenty years.’ I am in. Send me the engagement letter.”*

Sarah read the text twice. Then she set down her phone, put her hands flat on her desk, and let out a breath she felt like she had been holding since Friday.

She forwarded the text to David with one word: *“Reference.”* David’s reply came in thirty seconds: *“Linda told the truth. The truth worked. Candor. Imagine that.”*

Sarah sent the engagement letter to Allison within the hour. Three hundred vendor contracts. Forty-two thousand dollars. Seven business days. Their first engagement built entirely around the two-tier positioning strategy, won not by a perfect track record but by an honest one.

She thought about Mark Hensley, who had wanted to be second. Allison McLindon had just agreed to be first—not because Sarah had hidden the risks, but because Linda Torres had confirmed that when things went wrong, Candor fixed them transparently. The proof point Hensley needed was being built right now, one honest reference call at a time.



Sarah was tired but also wired. Allison McLindon’s engagement letter was already drafted—David had prepared the template and Sarah had customized the scope and pricing. Three hundred vendor contracts, fixed fee of forty-two thousand dollars, seven-business-day turnaround. Their first engagement designed from the ground up around the two-tier model.

She checked her phone. A text from Alex:

“How did Denver go?”

She typed back: *“She said yes. First client who was truly bought in on the full positioning story. Standard tier for the contract review, Advisory tier for the indemnification issues. One firm, two service levels, one relationship. And she mentioned eight hundred thousand in annual spend she wants to consolidate.”*

Alex’s reply came quickly: *“That is the template. Replicate it ten or twenty times and you have a Series A story.”*

Sarah put the phone away and watched the desert scroll past her car window—saguaros silhouetted against the city glow, the Superstition Mountains a dark mass on the eastern horizon. She did not feel triumphant. She felt the particular anxiety of someone who has just made a promise she is not entirely sure she can keep. Allison expected three hundred contracts analyzed in seven days. The team had never processed that volume under a real deadline with a real client waiting. If Candor buckled, if the quality slipped, if Priya’s retrieval layer choked on a document format they had not tested—Allison would not call back. And the eight hundred thousand dollars in annual spend would stay right where it was, split among four mediocre firms.

6.4 THREE HUNDRED CONTRACTS IN FIVE DAYS

The Allison McLindon engagement launched on a Wednesday morning. Priya had spent the previous four days configuring the healthcare compliance pipeline, a workflow she had prototyped using Claude Code in roughly a quarter of the time the original contract review pipeline had taken. The HHS business associate agreement requirements were loaded into the retrieval layer, cross-referenced against the current regulatory text. Priya had connected the pipeline

to two external regulatory databases—the HHS breach portal and a commercial healthcare compliance feed—so the system could pull current enforcement actions and regulatory guidance as it analyzed each contract.

“The pipeline will check each vendor agreement against the BAA requirements,” Priya said at the Monday standup. “But it will also pull recent HHS enforcement actions involving similar provisions. If the agency fined someone last quarter for exactly the kind of data handling gap we are seeing in one of Allison’s contracts, the system surfaces that context. The reviewer does not have to go looking for it.”

Three hundred contracts entered the pipeline Wednesday morning. By Thursday at noon, the AI had processed two hundred and seventy-eight of them, with the remaining twenty-two queued for a second pass due to document formatting issues the ingestion layer had flagged automatically. That validation step was one David had insisted on after the garbled-PDF problems from the Torres engagement.

On Thursday afternoon, the system flagged something unexpected. A cluster of forty-three vendor contracts, all with companies providing data analytics or population health services, contained non-standard data handling provisions that diverged sharply from the boilerplate. The AI had identified the pattern across the cluster: each contract permitted the vendor to retain de-identified patient data for “product improvement purposes” with no explicit time limitation and no audit rights for Allison’s company.

Individually, each provision looked like a minor deviation. In aggregate, the pattern created a systemic exposure. If HHS interpreted “product improvement” as a secondary use requiring explicit BAA authorization—and recent enforcement guidance suggested they might—Allison’s company had forty-three contracts that were arguably non-compliant. Not because of a missing signature or an expired term, but because of a subtle interaction between the data retention language and the BAA framework that individual contract review had difficulty surfacing.

Sarah’s senior reviewer on the engagement—a lawyer named Catherine with fourteen years of healthcare regulatory experience—validated the AI’s finding and added a layer of analysis the system had not reached. Catherine identified

that three of the forty-three contracts contained an additional wrinkle: they referenced a master services agreement from 2019 that pre-dated the current data handling provisions. The interaction between the legacy MSA terms and the newer addenda created an ambiguity about which data handling standard governed.

The AI had not flagged this because the 2019 MSAs were not in the document set; they existed only as references in the current contracts. Catherine recognized the issue from a similar structure she had seen in a previous role and escalated it to Sarah for Advisory-tier treatment.

David tracked the engagement metrics in real time from his operations dashboard. Average cycle time per contract: eleven minutes of AI processing plus eight minutes of human review, well under the fifteen-minute-per-contract budget they had planned. Confidence scores averaged ninety-one percent across the portfolio, with the data-handling cluster pulling the average down. Those contracts scored in the low eighties, which is exactly why they were routed to Catherine for deeper review.

Human override rate was fourteen percent: the reviewer modified the AI's analysis on roughly one in seven contracts. The override rate on standard provisions like governing law and term length was under three percent, while the rate on data handling and indemnification provisions was closer to thirty percent. That granularity would feed back into the platform, telling Priya's team exactly where to focus the next round of retrieval improvements.

The team delivered the complete assessment on Tuesday—day five of a seven-day commitment. The deliverable included the structured contract-by-contract analysis, the portfolio-level compliance dashboard, and a separate Advisory memorandum on the forty-three-contract data handling pattern, with Catherine's analysis of the legacy MSA issue and a recommended remediation strategy.

Allison's response came the following morning, in an email that Sarah read twice before allowing herself to feel anything about it.

"I shared your report with our CEO and our chief compliance officer last night. The data handling pattern you identified across the analytics vendors—we had no idea. Our previous firm reviewed a subset of these contracts last year and flagged nothing. Your system caught a systemic risk that individual reviews missed because

it could see the pattern across the full portfolio. The CEO's exact words were: 'This is what I have been asking legal to do for three years.' I am formally engaging you for the remediation work on those forty-three contracts under your Advisory tier. But more importantly—when can we talk about the rest of my spend? I have eight hundred thousand dollars a year which I am looking to shift, and right now I am questioning every one of my firm relationships."

Sarah forwarded the email to Alex, David, and Priya. She did not add commentary. The email spoke for itself. One engagement, delivered on time, under budget, with a finding that justified the entire AI-native thesis—not because the AI was smarter than a human lawyer, but because it could hold three hundred contracts in its analytical frame simultaneously and see patterns that sequential human review, no matter how diligent, would easily miss.

David replied first: *"Capture every metric from this engagement. This is our reference case."*

Priya replied second: *"Catherine's MSA catch is exactly the kind of thing we need to design for. The AI should have flagged the external document references and requested them. Adding that to the next sprint."*

Alex replied third: *"This is your Series A story. Do it again."*

One client who said yes did not make a business. It made a test. But it was the first test that felt like it mattered—the first engagement designed from the ground up around the strategy she had just spent weeks developing, not another opportunistic matter from her contact list. And she had told Allison the truth—that healthcare was new, that the first pass might not be perfect, that the firm was still proving itself. She had named the company Candor for a reason. Tonight, for the first time, the name felt like more than an aspiration.

The technology platform Priya was building could handle it. The five-layer architecture, the learning loops, the workflow orchestration were all taking shape. What Sarah had been missing was the strategic clarity to aim it all. Now she had that: mid-market GCs, the two-tier service model, the dual positioning in the matrix. Whether the execution would match the strategy was another question entirely.

She also knew—in the honest part of her mind that kept her up at night—that every strategic position carried risks. The Standard tier could commoditize:

if multiple AI-native firms pursued the same volume strategy with similar technology, price competition could erode margins until the position was no longer profitable. The Advisory tier required demonstrating a quality premium that clients might not be able to evaluate. The analysis might genuinely be better, but if the client could not see the difference, the pricing would not hold.

And running both tiers under one roof demanded operational discipline that most organizations lacked. If the volume mindset infected the advisory work, quality would suffer. If the advisory mindset infected the volume work, costs would rise. David's process engineering was not a luxury. It was the mechanism that kept the two systems from contaminating each other.

Strategy was a choice about where to compete and how to win, but it was not a guarantee of winning. It was a bet, informed by analysis and sharpened by the counsel of people like Alex and Rachel and David, but a bet nonetheless.

A strategy. A dwindling runway. A growing client list. Margins that needed to improve. And Allison McLindon, who had said yes when everyone else was still thinking about it. It was a start. Whether it was enough remained to be seen.

COMING NEXT WEEK

Chapter 7
The Org Chart No One Expected

Sarah has a strategy, a target market, a two-tier service model, and her first real client. What she does not have is an organization designed to deliver on any of it. Chapter 7 tears up the traditional law firm pyramid and rebuilds it for an AI-native world—new roles, new reporting lines, new economics of talent—and asks the question every founder eventually faces: who do you hire, in what order, and what do you do with the people who were hired for a firm that no longer exists?

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

ABOUT THE AUTHORS

Daniel Martin Katz, PhD, JD is Professor at Illinois Tech–Chicago Kent College of Law and Academic Director of the Bucerius Center for Legal Technology & Data Science. 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>