


# BUILDING AI-NATIVE PROFESSIONAL SERVICES FIRMS

*Strategy, Economics, and Execution*



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

## CHAPTER 8

# *Growing Pains: Talent and Culture in the AI-Native Firm*

CHAPTER 8 · 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
- ✓ Chapter 7 · The Org Chart No One Expected
- **Chapter 8 · Growing Pains: Talent and Culture** YOU ARE HERE
- 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 8: Growing Pains: Talent and Culture in the AI-Native Firm*

---

*Culture is not an initiative.  
Culture is the by-product of consistent behavior.*

— Seth Godin

### 8.1 SOMETHING IS GOING TO SLIP

The clients were coming. The pods were delivering. The revenue was climbing. But winning on the outside had created a different kind of problem on the inside.

Sarah had not slept well in three weeks. Not the anxious insomnia of her early months, when she would lie awake wondering if the whole venture was a mistake, but a different kind of wakefulness. The firm was growing, and the growth was breaking things.

She sat in her office at 9:40 a.m. on a Tuesday in late March (roughly nine months after the pod structure launched), the Phoenix sun already bright through the office windows, her laptop open to three separate dashboards. The first showed the engagement pipeline: thirteen active clients, twenty-two matters in various stages of completion, four new proposals awaiting client approval. The second showed the team: twelve people now, up from the nine who had

gathered for the all-hands at the pod launch the previous spring, spread across pods handling regulatory compliance, contract review, and due diligence. The third showed the numbers that Alex checked every Monday morning: \$2.5 million in run rate revenue, EBITDA margins hovering at 32 percent, cash runway of fifteen months.

The numbers were good. Not the hockey stick Alex had hoped for when he led the seed round, but real. Sustainable. Growing.

The people were another story.

Sarah scrolled to her email and reread the message that had kept her up the night before. It was from Megan Rivera, one of her three original lawyers, the ones who had been with the firm since the beginning.

*“Sarah, I need to talk. I’m drowning. I’m reviewing output on fourteen active matters. Joshua is on eleven. We can’t keep doing this. Something is going to slip, and when it does, it won’t be a small slip.”*

Sarah read the email twice. Megan was not someone who exaggerated. If she said she was drowning, the waterline was already above her chin.

That had been two months ago. They had tightened the quality processes, added structured checklists, implemented a two-reviewer protocol for any matter flagged above a certain risk threshold. But those fixes had only increased the load on the senior lawyers. Joshua Thornton (her first external hire, recruited in the early months), along with Megan Rivera and Ryan Gallagher, the two senior lawyers she had brought in to launch the pod structure the previous spring, were her three senior lawyers. All three were now reviewing everything, supervising everyone, and burning out.

Elena had left six weeks earlier. She had followed Sarah from the old firm on the strength of a conversation in a coffee shop and a promise that they would build something different together. She had been employee number two, had carried the legal workload with Joshua through the first months before the seed round closed, had floated between pods during the launch, had absorbed every crisis Sarah threw at her without complaint. When she asked for coffee at the end of a Friday afternoon, Sarah knew before they sat down.

“I’m not angry,” Elena had said. “I’m tired. My husband and I are trying for a second child. I need work I can finish at six. An in-house role at a healthcare company pays less but ends when the day ends. That matters to me right now.”

Sarah had not argued. She had asked the questions a founder asks when a key early employee leaves (what they needed, whether compensation or title would change the decision, whether they would take a sabbatical instead) and Elena had answered each one patiently. None of it was about Candor. It was about the life Elena wanted, and that life did not include being on call for every pod escalation at eleven on a Tuesday night. Sarah signed the paperwork the following Monday and walked Elena to her car when she packed up her desk. They still texted. But the firm was down its first believer, and the waterline on Joshua, Megan, and Ryan had risen accordingly.

Sarah closed her laptop and poured a second cup of coffee. She needed to hire. The question was: who?



David Park arrived at the office at 8:15, as he always did, carrying a reusable coffee mug and a notebook thick with Post-it flags. David had been the firm’s fourth hire, brought on as COO after Alex Hawthorne introduced them, convinced that Sarah needed someone who understood production systems. His background in healthcare operations consulting, applying Lean and Six Sigma to clinical workflows, had translated to legal services better than anyone expected. He had mapped the firm’s workflows, identified bottlenecks, built the pod structure, and designed the quality dashboards. He spoke a language the lawyers found foreign: cycle time, throughput, defect rate, standard work. But the concepts translated.

This morning, David walked into Sarah’s office without knocking. He had a habit of reading her mood from the state of her desk. Today, the desk was clean except for the laptop and the coffee, which meant she was thinking hard about something unpleasant.

“Megan’s email?” he asked.

“You read it too?”

“She copied me.” He sat down across from her. “She is right. The reviewers are the constraint. We can push more work through the AI pipeline, but it all stacks up at the human verification stage. It is a classic bottleneck.”

“I know what it is. I need to know what to do about it.”

David opened his notebook. “Hire. But hire differently than you have been hiring.”

“Meaning?”

“Your three senior lawyers are excellent. They came from traditional firms. They had years of experience doing the work before they learned to supervise AI doing the work. That experience is valuable: they know what good looks like because they have produced it themselves.” He paused. “But that background also means they had to unlearn habits. Megan still catches herself wanting to rewrite AI-generated summaries from scratch. Joshua argues with the checklists because he trusts his judgment more than any process. Ryan is the best of the three at working within the system, but even he spends time on tasks the AI handles well because he cannot quite let go.”

Sarah knew all of this. She lived with it every day. But she and David had underestimated the degree of unlearning that was required for more experienced lawyers to work differently. “Old dogs, new tricks or whatever the saying is...”

“So what’s the alternative?”

“Expand your Legal Quality Analyst bench, but hire JDs this time.” David leaned forward. “Karen and Wei are excellent. Non-legal backgrounds, strong analytical skills, exactly what we needed to get the pods running. But you have a bigger problem than verification capacity. You may eventually have a senior lawyer pipeline problem. Alex told you in the board meeting: the constraint on growth is finding experienced lawyers willing to work in this model. So build your own. And yeah also the culture here needs folks who do not have to do so much unlearning.”

Sarah saw where he was going. “Hire junior lawyers into the Legal Quality Analyst role. They learn the system from the quality side: verifying, correcting, understanding error patterns. And the best ones grow into Senior Lawyers who already trust the system because they grew up inside it.”

“Exactly. In healthcare operations, when we introduced automated quality checks for clinical workflows, the best reviewers were not the ones who had done manual chart reviews for twenty years. They were the ones we trained from scratch on the new system. The veterans were better at catching certain things, but they were slower, and they fought the process. The new hires did not know any other way.” He paused. “You are not hiring associates. You are building a bench. Eighteen to twenty-four months in the Legal Quality Analyst role, learning every workflow, every error pattern, every client context. Then they step up into Senior Lawyer roles in new pods—and they do not have to unlearn anything.”

“You want me to hire lawyers who don’t know how to practice law the old way.”

“I want you to hire lawyers who will learn to practice law the way your firm practices it. There is a difference.”

## 8.2 CULTIVATING YOUR OWN SEED CORN

### JUNIOR LEGAL QUALITY ANALYST (JD TRACK)

The Legal Quality Analyst role, established in the pod model, verifies AI-generated legal analysis, identifies systematic errors, and feeds corrections into the learning pipeline. The firm’s first LQAs (Karen and Wei) came from non-legal backgrounds (pharmaceutical QA and data analytics). The JD track expands the role to junior lawyers with technical aptitude, creating a pipeline: Junior LQA → Senior LQA → Senior Lawyer. Requires a JD and bar admission plus demonstrated comfort with AI systems and data analysis. Does not require BigLaw experience, class rank, or law review credentials. The bench that solves the senior lawyer pipeline problem.

Sarah posted two positions the following week: Junior Legal Quality Analyst. The job description was unlike anything she had seen in legal hiring. It

required a JD and bar admission. It also required demonstrated technical aptitude: comfort with AI systems, experience with data analysis, willingness to learn workflow engineering. It explicitly did not require BigLaw experience. It explicitly did not require a class rank or a law review note. The marginal preference for JDs was deliberate: Karen and Wei had proven that non-lawyers could excel in the role, but lawyers who learned the system from the quality side would eventually become the Senior Lawyers the firm needed to scale.

The applications were revealing. She received forty-seven resumes in the first week. Most came from recent graduates or lawyers with one to two years of experience at small firms or in-house roles. Some had computer science minors, like Sarah herself. Others had worked in legal technology before law school. A few had non-traditional backgrounds entirely: one applicant had been a quality engineer at a semiconductor company before going to law school at thirty.

Sarah interviewed twelve candidates. She hired two: Anika Sharma and Nolan Webb.

Anika had graduated from a law school in California eighteen months earlier and had followed her husband to Phoenix, where he was undertaking a medical residency at a local hospital. She had passed the bar on her first attempt but had struggled to find a traditional associate position in Arizona. She was new to the market, had no local network, and the mid-market firms that interviewed her wanted someone who already knew how to manage a caseload. What she did have was a computer science degree from before law school, two years of experience as a data analyst, and an almost unsettling comfort with AI systems. During her interview, Sarah had asked her to review an AI-generated contract summary and identify errors. Anika found three issues in twelve minutes—faster than Megan typically managed.

Nolan had a more unusual path. He had practiced for two years at a small plaintiff's firm in Tucson before deciding that traditional litigation was not for him. He had spent the past year teaching himself prompt engineering and building a side project that used AI to analyze medical records for personal injury cases. His legal experience was thin. His technical skills were strong. And his answer to Sarah's final interview question—"What do you think lawyers will be doing in ten years?"—had been the best she heard: "Lawyers will be doing what

they should have always been doing. Thinking. Judging. Advising. Everything else will be handled by systems that do it better than we ever did.”

The reaction from the existing team was immediate.

Megan was cautiously supportive. “We need the help. I don’t care if they come from law school, BigLaw or a lemonade stand, as long as they can review work accurately.”

Ryan was pragmatic. “Train them well, give them structure, and they’ll be fine. Most of what my old firm taught me was irrelevant to what I do here anyway.”

Joshua was not supportive.



Joshua Thornton had been Sarah’s first hire. She had recruited him from a well-regarded Phoenix firm where he had spent six years building a reputation as a meticulous regulatory lawyer. He was smart, thorough, and fast. He had left his old firm because he believed in the AI-native thesis, because he wanted to build something new, and Sarah suspected, in part, because he liked the idea of being important at a small firm rather than anonymous at a large one.

For the first nine months, Joshua had been essential. He had built the regulatory compliance workflow from scratch, trained the AI agents on Arizona and federal regulatory patterns, and served as the firm’s quality backstop on their most complex matters. When the Torres engagement hit (350 contracts in two weeks, the firm’s first true stress test), it was Joshua who carried the heaviest load, reviewing more contracts than anyone else on the team, working nights to meet the deadline. His expertise had caught errors the AI missed, and his judgment calls on complex provisions had been the difference between a credible deliverable and an embarrassing one.

But something had shifted in Joshua during those months. The Torres engagement had confirmed his belief that the firm’s quality depended, in the end, on experienced lawyers exercising independent judgment. He saw the checklists and structured processes that David introduced not as improvements but as

insults: bureaucratic constraints imposed by a non-lawyer on the professional discretion that defined good legal work.

The tension surfaced in small ways at first. Joshua would skip pod meetings because he “had real work to do.” He would complete quality reviews but annotate the checklists with curt notes: “This step is redundant,” or “Process adds twenty minutes for no value.” He referred to David’s systems as “the factory” with an inflection that made clear it was not a compliment.

The hiring of Anika and Nolan brought the tension to the surface.

It happened on a Wednesday, in the small conference room the firm used for team meetings. David was presenting the onboarding plan he had designed for the two new hires—a structured twelve-week program modeled on the clinical training rotations he had used in healthcare operations.

“Week one through three, they shadow the AI pipeline team,” David said, projecting his slides. “They learn how the agents process documents, how the risk classification works, how the outputs are structured. No client work yet. Just learning the system.”

“Weeks four through six, they rotate through each pod. Two days in regulatory, two days in contract review, two days in due diligence. They observe, they ask questions, they start doing supervised practice reviews on completed matters where we already know the right answers.”

“Weeks seven through nine, they start live reviews under direct supervision. Every review is checked by a senior lawyer. Every discrepancy is documented, discussed, and used as a training case.”

“Weeks ten through twelve, they begin handling reviews independently but with random audits. We track accuracy rates, cycle times, and flag rates. They do not graduate to solo work until they hit 97 percent accuracy over a two-week window.”

 CANDOR'S 12-WEEK CERTIFICATION PROGRAM

**Weeks 1–3:** Shadow the AI pipeline team. Learn document processing, risk classification, output structure. No client work.

**Weeks 4–6:** Rotate through each pod. Observe, ask questions, practice on completed matters with known answers.

**Weeks 7–9:** Live reviews under direct supervision. Every discrepancy documented and used as training.

**Weeks 10–12:** Independent reviews with random audits. Must achieve 97% accuracy over a two-week window to graduate.

Practical assessments use real (anonymized) files with deliberately introduced errors, the same error types encountered in actual engagements.

Megan nodded. Ryan was taking notes. Sarah watched Joshua, who was leaning back in his chair with his arms crossed.

“Questions?” David asked.

“I have one,” Joshua said. His voice was measured, but Sarah could hear the edge. “These two people have zero experience in actual legal practice. Anika has never handled a live client matter. Nolan spent two years at a plaintiff’s shop doing intake work. And we’re going to put them in front of clients after twelve weeks of what amounts to an assembly-line orientation?”

David started to respond, but Joshua continued.

“I spent six years learning regulatory law. Six years reading cases, drafting memos, sitting in on hearings, getting torn apart by partners who expected perfection. That’s how you develop professional judgment. You don’t develop it by checking boxes on a quality audit form.”

The room went quiet. Sarah let the silence hold for a moment before speaking.

“Joshua, your experience is exactly why you’re valuable here. Nobody is saying otherwise. But Anika and Nolan aren’t being hired to replace you. They’re joining as Junior Legal Quality Analysts, the same role Karen and Wei hold, but

with JDs. They learn the system from the quality side. They verify, they correct, they build judgment inside our workflows. And in eighteen months, two years perhaps even five years, the best ones step up into Senior Lawyer roles in new pods. This is how we solve the growth constraint Alex flagged.”

“They won’t learn law at all,” Joshua said. “They’ll learn to operate a machine.”

“They’ll learn to practice law the way this firm practices it,” Sarah replied. “With AI handling the first pass and humans providing the judgment. That’s not less rigorous. It’s differently rigorous. And they won’t have to unlearn anything. They’ll learn our system from scratch, which means they’ll trust it in ways that—honestly—some of us still struggle with. This is how we build the bench.”

Joshua looked at her for a long moment. Then he picked up his notebook and left the conference room without another word.

### 8.3 DRAFT AND DEVELOP

In professional sports, there are generally two ways to build a roster. You can chase free agents, paying a premium for proven talent who arrive with established habits and established price tags. Or you can draft and develop, investing in raw potential and shaping it inside your system. The sports franchises that sustained success optimized as between marquee free agent signings and development of internally cultivated talent.

Sarah was building a draft-and-develop firm. The senior lawyers she had recruited (Megan, Ryan, Joshua) were the free agents, proven talent who could contribute immediately but came with habits formed elsewhere. Anika and Nolan were the draft picks, raw but moldable, learning the system from the inside out. The twelve-week certification program was the training camp.

The onboarding program David designed reflected a set of principles that any good talent development system requires. But living through it was messier than any roadmap suggested.

The first principle was assessment before action. Before hiring Anika and Nolan, Sarah and David had mapped the firm’s existing capabilities against its

BUILDING THE ROSTER: TWO APPROACHES		
	FREE AGENT	DRAFT & DEVELOP
Cost	High: market-rate salary	Lower: junior salary + equity
Ramp Time	Fast but requires unlearning	12 weeks; no unlearning needed
Cultural Fit	Variable: habits from elsewhere	High: shaped from day one
System Trust	Earned slowly; skepticism common	Native: they know no other way
Retention Risk	Higher: portable skills, options	Lower: equity + professional identity
Bench Building	No: fills today's gap only	Yes: solves tomorrow's pipeline
<b>VERDICT: Need both. Shifting toward draft-and-develop as the firm scales.</b>		

**Figure 8.1:** Two approaches to building the roster. Candor needs both free agents and draft-and-develop talent, but the ratio shifts toward draft-and-develop as the firm scales and the bench deepens.

needs: strong domain expertise in regulatory and contract law, adequate AI supervision skills among the senior team, weak process discipline, and nonexistent bench depth for the verification layer that was the firm’s core value proposition. The Torres engagement’s 18 percent error rate and chronic corner-cutting under deadline pressure had made the process gap painfully clear.

The gap was specific and actionable. They did not need more senior lawyers. They needed a middle layer of professionals who could handle routine AI supervision reliably, freeing the senior team to focus on complex matters and quality exceptions.

The second principle was structured development over apprenticeship. Traditional law firms develop talent through osmosis: junior lawyers absorb skills by working alongside senior ones, with feedback that ranges from thorough to nonexistent depending on the partner’s temperament. David’s twelve-week program replaced osmosis with intentional design. Every skill had a training module. Every competency had a measurable standard. Every transition, from observer to supervised practitioner to independent reviewer, had clear criteria.

This was the operations mindset David brought to the firm. In his health-care consulting years, no one was placed on a clinical workflow until they had demonstrated competency at each station through a structured rotation. The idea that a junior associate would be given a pile of documents and told to “figure it out” struck him as negligent. “In healthcare operations, that is how patients get hurt,” he told Sarah. “In law, that is how clients get hurt.”

The third principle was certification on quality standards. David’s program required Anika and Nolan to pass assessments at each stage: not perfunctory quizzes, but practical evaluations using real (anonymized) matter files. They reviewed AI-generated outputs where David had deliberately introduced errors: the same types the firm had encountered in actual engagements, including the force majeure cross-referencing failure from the logistics engagement. They had to identify the errors, explain why the AI had made them, and describe what the correct output should look like.

David called it the flight simulator. Commercial pilots did not learn to handle engine failures by waiting for one to happen at thirty thousand feet. They trained in simulators that reproduced every failure mode the aircraft could experience, over and over, until the correct response was muscle memory. David built the legal equivalent. He curated a library of hundreds of AI outputs, each seeded with errors drawn from the firm’s actual engagement history: missed clause interactions, misclassified risk ratings, outdated regulatory references, jurisdiction-specific provisions the model had treated as generic. The difficulty scaled with the trainee’s progress. Early simulations contained obvious errors, the kind a careful reader would catch. Later ones contained the subtle failures that had caused real problems on real engagements, the kind that required not just attention but judgment.

“Anyone can catch a typo,” David told Anika during her week-eight assessment. “I need you to catch the thing that looks right but is not. The AI is very good at producing output that looks right. Your job is to know the difference.”

Priya automated the simulator so it generated fresh scenarios weekly, pulling from new engagement data and recent knowledge base updates. The system tested what the trainees had learned and, just as importantly, revealed what the AI had not. Every simulation failure that stumped a trainee became a diagnostic:

was the gap in the trainee's knowledge, or in the platform's? Often it was both. The flight simulator trained the people and improved the system simultaneously.

The assessments were rigorous enough that Megan, watching Anika complete one during week eight, admitted she was not sure she would have passed it on her first attempt.

### 8.3.1 *The Skills Taxonomy at Work*

David's onboarding program organized skills into three tiers: foundation skills that everyone needed, role-specific skills tied to particular functions, and leadership skills for those driving the firm's evolution. Each tier had concrete training content.

The foundation skills came first. AI tool proficiency meant not just knowing which buttons to press but understanding how the firm's AI pipeline processed documents, why certain document types produced higher error rates than others, and how to construct prompts that elicited more accurate outputs from the system. Critical evaluation meant developing a systematic approach to reviewing AI output: not reading every word the way a traditional associate would read a brief, but scanning for the specific failure patterns that the firm's quality data showed were most common.

Process execution meant following the workflows David had documented, including the ones that felt tedious and redundant. Those workflows existed for reasons that the Torres engagement's error rates, cross-referencing gaps, and near-misses during rapid growth had made viscerally clear.

Role-specific skills developed during the pod rotations. In the regulatory pod, Anika learned to evaluate AI-flagged compliance issues against a decision tree that the pod had developed over dozens of engagements. In the contract review pod, Nolan learned to calibrate his judgment against the firm's historical accuracy data. If he was flagging risks at a rate well above or below the baseline, something in his evaluation process needed adjustment.

Leadership skills would come later. For now, the goal was competence, then consistency, then the kind of pattern recognition that only develops through volume and feedback.



The first month was harder than Sarah expected.

Anika adapted quickly. Her data analysis background gave her an intuitive sense for patterns in AI output: she could look at a batch of flagged contracts and identify systemic issues before the senior reviewers had finished their first file. By week six, she was catching errors at a rate that matched Ryan's, though her domain knowledge was still developing.

#### ANIKA VS. NOLAN: TWO PATHS TO AI-NATIVE PRACTICE

**Anika Sharma:** CS degree + data analyst background, California law school. Adapted quickly; her pattern recognition skills let her scan AI output batches and spot systematic issues before senior reviewers finished their first file. By week six, catching errors at a rate matching senior staff.

**Nolan Webb:** Two years plaintiff's litigation in Tucson, self-taught prompt engineering. Struggled initially; wanted to engage with substance rather than verify outputs. Breakthrough in week six when he caught a non-compete conflict with a recently amended statute that the AI missed entirely, leading to a knowledge base update.

Different backgrounds, different learning curves, both essential. The system benefits from cognitive diversity.

Nolan struggled initially. He had strong opinions about what “practicing law” meant, and the structured rotation challenged those opinions. He wanted to engage with the substance of each matter: read the contracts himself, form independent views, argue for interpretations that the AI hadn't considered. David had to remind him, more than once, that his job was to verify AI output, not to redo the AI's work.

“I know you want to engage with the material,” David told Nolan during a coaching session in week four. “And you will. But first, you need to learn the

system. Trust the process. The judgment comes after you have mastered the mechanics.”

Nolan resisted for another week. Then, during a review of a batch of employment agreements, he caught an issue that the AI had missed entirely: a non-compete provision that conflicted with a recently amended Arizona statute. The AI had flagged the provision as standard. Nolan recognized the conflict because he had spent his Tucson years handling employment disputes and knew the statute had changed.

The catch was significant. Sarah flagged it for the team and used it as a training example. But the lesson was not just that Nolan had good legal instincts. The lesson was that the system had a gap: the AI’s training data did not yet reflect the recent statutory amendment, and Nolan’s catch led to a knowledge base update that would prevent the same miss on future matters. It also surfaced a broader need. “We need a statutory and regulatory surveillance system,” Priya said at the next standup. “It is the digital equivalent of the old school ‘pocket parts’ my professor used to talk about. The law changes. Our knowledge base has to change with it, automatically, not whenever someone happens to catch it.” She added it to the platform roadmap that week.

“That is the flywheel,” David said when Sarah described the incident at their weekly operations meeting. “Every error we catch, every gap we fill, makes the system better. The humans make the AI smarter. The AI makes the humans faster. The whole thing compounds.”

Sarah nodded. She had described this concept to investors a dozen times. Hearing David articulate it from an operations perspective, without any of the pitch-deck language, made it feel more real.

The incident also crystallized something she had been sensing but had not yet articulated: experienced lawyers were not displaced by AI. They were elevated by it. Nolan’s two years of employment litigation, the pattern recognition he had developed reviewing hundreds of employment agreements, the instinct for which provisions actually mattered in practice, the knowledge of recent statutory changes that came from working in the field: none of that was replicated by the AI. What the AI did was handle the routine analysis that had consumed most of a junior lawyer’s time, freeing the judgment to focus where it mattered most.

A new graduate with no employment law experience could verify AI output against the system's checklists competently enough. Only Nolan could have caught the non-compete conflict, because catching it required knowing something the system did not yet know. In the AI-native firm, experience was not a depreciating asset. It was the asset the technology made more valuable. The lawyers who had spent years developing domain expertise were sitting on exactly the resource that the new model rewarded.



The conflict with Joshua did not resolve.

It worsened. Over the weeks following the team meeting, Joshua became increasingly isolated. He completed his work (his output was as good as ever, his quality reviews thorough and precise) but he disengaged from the firm's culture.

He stopped attending the weekly retrospectives that David had instituted, where the team reviewed quality metrics, discussed near-misses, and brainstormed process improvements. He stopped mentoring Nolan, who had been assigned to shadow him for the regulatory pod rotation. He started taking calls in his car during lunch, and Sarah suspected he was talking to recruiters.

The breaking point came in early August, on a Tuesday morning when the air outside was already dry and fierce before eight o'clock.

Sarah had introduced a revised compensation framework the previous week. The scaffolding Sarah had put in place at the pod launch (base salary, pod-level outcome bonus, and equity vesting over four years) had worked, but it had been layered on top of a quarterly individual accelerator tied to review accuracy and client satisfaction scores, a holdover she had kept for the senior lawyers who had joined before the pods existed. The revision removed the individual accelerator entirely. All variable compensation would now flow through the pod: shared metrics on client retention, quality scores, turnaround time, and margin contribution, distributed equally among pod members with seniority adjustments. The equity component was unchanged. The Junior Legal Quality Analysts would participate in the same pod bonus pool from the day they graduated the certification program.

The change was deliberate. Sarah had studied the compensation models described in the analytical literature on AI-native firms: outcome-based rather than input-based, team-oriented rather than individual, with equity participation that aligned long-term interests. She had discussed the design with Alex, who had seen similar structures work in his portfolio companies. “You’re building a team sport,” Alex had told her. “Pay like it’s a team sport.”

Megan and Ryan accepted the change without enthusiasm but without resistance. Megan’s concern was practical: “As long as the pod targets are things I can actually influence, I’m fine. If I get penalized because the AI pipeline has a bad week, that’s not fair.” Sarah had built in adjustments for system-level issues.

Joshua’s reaction was different.

He came to Sarah’s office at 7:30, before anyone else had arrived, and closed the door behind him.

“I can’t do this,” he said.

### THE JOSHUA THORNTON PROBLEM

Talented professionals who built their identity around individual expertise and autonomous judgment may not thrive in process-driven, team-accountable AI-native environments. The resistance is not about competence; Joshua was one of the firm’s strongest performers. It is about identity: when structured verification replaces individual discretion as the quality mechanism, some professionals experience it as a demotion in the nature of the work itself. Not every professional will adapt, and the honest response is to acknowledge this rather than pretend otherwise.

“Can’t do what, specifically?”

“Killing the individual bonus. The checklists. The twelve-week training program for people who’ve never handled a real matter. The quality audits and the retrospectives and the certification assessments.” He sat down heavily. “I didn’t leave my old firm to join a factory, Sarah. I left to practice law in a better way. And I thought that’s what you were building.”

“I am building a better way to practice law.”

“No. You’re building a better way to process legal work. That’s not the same thing.” Joshua leaned forward. “Practicing law means exercising professional judgment. It means knowing the right answer because you’ve seen enough wrong answers. It means having the autonomy to make calls that a checklist can’t capture. Every process David adds, every metric he tracks, every form he makes us fill out. It erodes that. It turns lawyers into operators.”

Sarah had heard versions of this argument for months. She had been patient because Joshua was talented and because he was partly right. Professional judgment did matter, and processes could become rigid. But she was also tired, and she had spent too many nights haunted by Megan’s email—*something is going to slip, and when it does, it won’t be a small slip*—to accept the premise that professional judgment alone was sufficient.

“Joshua, I’ve watched the pattern since the Torres engagement. You approved flagged deviations without explanation. Your shorthand produced client-visible errors. And now you’re disengaged: skipping mentoring sessions, refusing to attend retrospectives. The process isn’t eroding your judgment. Your resistance to the process is eroding the team.”

The words landed like a closing argument. Joshua’s jaw tightened.

“That’s not fair. I was overloaded. We all were.”

“I know. That’s exactly my point. Professional judgment fails when people are overextended, when there’s no system to catch what individuals miss, when quality depends entirely on how much sleep someone got the night before.” Sarah kept her voice steady. “The processes David built aren’t there to replace your judgment. They’re there to catch the moments when your judgment—anyone’s judgment—isn’t enough.”

Joshua was quiet for a long time.

“I’ve been talking to Morrison and Associates,” he said finally. “They’ve offered me a senior associate position. Partnership track. Traditional firm, traditional practice.”

Sarah had known this was coming. She had known since the day he walked out of the conference room. But knowing did not make it easier.

“Is that what you want?”

“I want to practice law the way I was trained to practice law. I don’t think that’s wrong.”

“It’s not wrong. It’s just not what we’re building here.”

They talked for another twenty minutes. Sarah did not try to convince him to stay. She told him she respected his decision, that his work had been essential to the firm’s survival, that the regulatory workflow he had built would carry his fingerprints for years. She offered a generous severance: three months of salary and acceleration of his first year of vested equity. Joshua accepted.

He was gone by the end of the week.

## 8.4 THE DOOR CLOSSES BEHIND JOSHUA

Joshua’s departure forced Sarah to confront something she had been avoiding: the firm’s culture was not accidental. It was being forged through conflict, through loss, through the daily accumulation of choices about what to tolerate and what to demand.

She sat with it that evening and tried to be honest about what had gone wrong. The firm’s readiness for change was not uniform. Some dimensions were strong: she had full authority as founder, the technology platform was purpose-built, and the investors were patient. But culture was the weak link because the culture did not yet exist when Joshua joined. It was still being negotiated among people with fundamentally different assumptions about what good work looked like.

Joshua had been exactly what the firm needed in terms of skill. He was not what the firm needed in terms of values. The gap between the two had widened until it became unbridgeable. He was not resistant to all change. He had left a traditional firm, after all. He was resistant to the specific type of change the firm required: a shift from individual autonomy to collaborative process, from professional discretion to structured verification, from “I know best” to “the system helps us all be better.”

Sarah thought about what had driven the resistance. It was not fear of obsolescence. Joshua was confident in his abilities. It was not a concern about competence. He was one of the firm's strongest performers. It was about identity. He had built his professional self-image around being a skilled practitioner who exercised independent judgment. The firm's evolution toward process-driven quality felt like a demotion, not in title or compensation, but in the nature of the work itself. And that was something no compensation package or career path could fix.

Sarah could not have prevented this. She could have managed it differently: earlier conversations, more explicit discussions about the firm's direction, perhaps a role that gave Joshua more autonomy on complex matters while still requiring compliance with quality processes on routine work. But she doubted the outcome would have been different. Some professionals will not thrive in AI-native environments, and the honest response is to acknowledge this rather than pretend that everyone will adapt.

What mattered now was what the departure signaled to everyone else.



Sarah called an all-hands meeting the morning after Joshua left. All thirteen remaining team members gathered in the conference room: the lawyers, the engineers, David, the operations coordinator. Sarah stood at the front without slides, without notes.

“Joshua left yesterday. Most of you know this already.” She paused. “I want to be honest about what happened and what it means.”

“Joshua is an excellent lawyer. His work here was outstanding, and the regulatory workflow he built is one of the things that makes this firm good at what we do. He left because the direction we're heading—more process, more structure, more team-based accountability—isn't the kind of practice he wants. And I respect that.”

She looked around the room. Anika was attentive. Nolan looked uncomfortable. Megan was watching Sarah with the careful expression of someone evaluating whether leadership was going to be honest or performative.

“Here’s what I want you to take from this. We are building something that doesn’t have a template. There is no established model for how an AI-native law firm works, how it develops talent, how it defines quality, how it compensates people. We are making it up as we go—thoughtfully, I hope, but we are making it up.”

“That means not everyone who starts this journey will finish it. Some people will decide this isn’t for them. That’s okay. What isn’t okay is pretending we can build this firm without changing how we work. The processes David builds are not optional. The quality standards are not suggestions. The team-based compensation is not an experiment we might reverse. These are the foundations of what we’re becoming.”

She took a breath.

“Three values. I want to name them because I think they’ve been implicit and they need to be explicit.”

“First: experimentation over perfection. We are going to try things that don’t work. New workflows, new tools, new ways of reviewing AI output. When something fails, we learn from it and adjust. We don’t punish people for experiments that go wrong. We punish people for refusing to experiment.”

“Second: outcomes over inputs. I don’t care how many hours you work. I care about whether the client got excellent work product, on time, at the quality standard we’ve committed to. If you can do that in six hours, don’t sit at your desk for ten to look busy. If it takes twelve hours because the matter is complex, we’ll talk about how to make it more efficient next time.”

“Third: team over individual. The pod structure exists because no individual—no matter how talented—can consistently deliver the quality our clients need. The system catches what individuals miss. The team compensates for each person’s blind spots. When we win, we win together. When something goes wrong, we fix it together.”

The room was quiet. Then Megan spoke.

“I have a question. Does this mean we’re going to stop hiring experienced lawyers?”

“No. We need experienced lawyers. We need people like you and Ryan who can handle complex judgment calls and mentor the junior team. But we’re also going to hire people like Anika and Nolan, people who learn our system from the inside rather than bringing habits from a different system. The mix matters.”

David added a thought. “In healthcare operations, the best teams have a blend of experienced clinicians who understand the work deeply and newer staff who are fluent in the current quality systems. Neither group alone is as effective as the combination.”

Megan considered this. “Fine. But if I’m going to mentor Anika and Nolan on judgment calls, I need fewer matters on my plate. I can’t develop people and review fourteen active matters at the same time.”

“Agreed,” Sarah said. “That’s the next thing David and I are working on: restructuring the pod assignments so that the senior team has capacity for mentorship and complex work, and the junior team is handling more of the routine verification.”

The meeting ended. People filed out. Nolan stopped at the door and turned back.

“For what it’s worth,” he said, “I think this is right. The structure, the training, the values. It’s why I came here.”

Sarah smiled. “Good. Now go prove it.”

## 8.5 CULTURE FORGED, NOT JUST DECLARED

Culture, Sarah was learning, does not arrive through declarations. It arrives through repetition. The three values she articulated in the all-hands meeting (experimentation over perfection, outcomes over inputs, team over individual) became real not because she said them but because she and David reinforced them through daily decisions.

When Anika proposed a modification to the contract review workflow that would skip a verification step on low-risk provisions, Sarah did not dismiss the idea. She asked Anika to run a controlled test: process fifty contracts with the modification and fifty without, then compare error rates. The modification

reduced cycle time by 18 percent with no measurable increase in errors. It was adopted firm-wide. Anika received a team-wide acknowledgment at the next retrospective. That was experimentation over perfection in practice.

When Ryan finished a complex regulatory analysis in seven hours that the pod had estimated would take twelve, Sarah did not ask him to find more work to fill his day. She asked him to document what he had learned so the pod could update its estimation model. Then she told him to go home early. That was outcomes over inputs in practice.

When a quality audit revealed that one pod was consistently flagging more issues than the others (not because the work was riskier, but because one reviewer was applying an overly conservative threshold), the team discussed it openly at the retrospective. They did not single out the individual. They adjusted the threshold guidance for everyone and ran a calibration session using sample matters. That was team over individual in practice.

These small moments accumulated. By October, three months after Joshua's departure, the culture had shifted in ways that were visible even to clients. Allison McLindon at Summit Health Partners, who had become a steady client after her pilot engagement, told Sarah during a quarterly review: "Your team feels different from other firms. More collaborative. More transparent about how the work gets done. I appreciate that."

### *8.5.1 Compensation as Culture Signal*

The compensation restructuring that had been the proximate cause of Joshua's departure became one of the firm's most important cultural signals. The pod-level outcome metrics created a shared accountability that reshaped daily behavior.

Under the old individual-bonus structure, there had been subtle competition among the lawyers: who reviewed more contracts, who caught more errors, who maintained higher accuracy scores. The competition was not destructive, but it was not collaborative either. Lawyers would occasionally sit on a useful insight rather than share it immediately, wanting credit for the catch rather than the process improvement.

Under the new pod structure, the incentives reversed. If Anika discovered a pattern in AI errors that could help the whole pod, sharing it immediately improved the pod's quality score, which improved everyone's bonus. If Nolan was struggling with a particular type of review, the pod had a financial interest in helping him improve rather than letting him flounder. Knowledge sharing became economically rational, not just culturally encouraged.

The equity participation added a longer time horizon. Every team member now had a stake in the firm's overall success: not just their pod's quarterly metrics but the firm's growth, its margin improvement, its client retention. When Alex visited for a board meeting in August and asked the team what they were most focused on, Megan's answer surprised him: "The data flywheel. Every engagement makes the AI better, which makes us faster, which lets us serve more clients, which gives us more data. We're building compound interest."

Alex had looked at Sarah. "She sounds like a founder."

"That's the point."

The compensation model was not without challenges. Calculating pod-level bonuses without the individual accelerator layer required more administrative work than the prior hybrid scheme. Some team members found it unsettling to have their compensation partially dependent on colleagues' performance. And the equity vesting schedule meant that the full alignment effect would take years to materialize.

But the signal was clear: this firm valued collaboration, outcomes, and long-term thinking. People who thrived on individual competition and short-term recognition would not be comfortable here. That was intentional. Sarah had learned from the Joshua experience that cultural fit was not a secondary consideration in hiring. It was primary.

## 8.6 BUILDING THE TALENT PIPELINE

By September, the firm had stabilized. Anika had passed her twelve-week certification and was handling independent reviews on contract matters. Nolan was two weeks behind her, still in the supervised review phase for regulatory work, but his accuracy had reached 96 percent and was trending upward. Megan and

Ryan reported that their workload had decreased from unsustainable to merely heavy. A meaningful improvement.

But Sarah knew that stability was temporary. The pipeline showed increasing client interest. Two large potential engagements, a multi-state regulatory review and a series of commercial contract analyses for a real estate investment trust, would require more capacity than the firm currently had. She needed to think about the next wave of hiring.

David had been thinking about it too. He presented Sarah with a framework he called the “talent blueprint,” a document that mapped the firm’s growth targets against the capabilities required at each stage.

“At twelve people and \$2.5 million in run rate revenue, we need what we have: three senior lawyers for complex work and mentorship, two junior lawyers for routine verification, a technical team for the platform, and an operations staff for everything else.”

“At twenty people and \$5.5 million in run rate revenue, which is your twelve-month target, we need to add capacity in a specific sequence. First, two more junior lawyers trained on the same twelve-week program. Second, one mid-level lawyer who can bridge between the senior team and the juniors, someone with enough experience to handle moderately complex matters independently but enough flexibility to work within the system. Third, one more engineer to handle the platform improvements that our quality data is generating.”

“The key insight is the ratio,” David continued. “In a traditional firm, the ratio of senior to junior lawyers might be one to four (and sometimes more). In our model, it’s closer to one to two, because the AI handles the work that juniors would have done. But the seniors need to spend more time on supervision and mentorship, which means their own matter load has to be lower. The math is different.”

Sarah studied the blueprint. The logic was clear: selective retention of high performers, targeted hiring for specific capability gaps, and advancement of professionals who demonstrated AI-native competence. The firm was compressing what would typically be a twelve-to-twenty-four-month optimization phase into something closer to six months, the advantage of building AI-native from scratch rather than transforming an existing organization.

“What about the mid-level hire?” Sarah asked. “That’s the hardest one. Someone with enough experience to bring real judgment but enough openness to work within our system.”

David nodded. “I’d look for someone who’s already frustrated with traditional practice. Not someone who’s fleeing—that’s how you get another Joshua. Someone who’s actively looking for a different way to work. Maybe from an ALSP, or an in-house legal operations role. Someone who’s seen what process can do for legal work and wants more of it, not less.”

“A unicorn.”

“A rare hire. But they’re out there. The legal profession is producing more people who think like this than it used to. Remember there is someone just like you out there sitting in a meeting room with the idea that there must be a better way. We just have to find them.”



The data flywheel that Megan had described to Alex was, by October, producing measurable results.

The firm tracked a metric David had invented: “system learning rate,” defined as the percentage improvement in AI accuracy per hundred engagements processed. In the firm’s first six months of operation, the learning rate had been high: each batch of engagements revealed significant gaps in the AI’s training data, and the corrections produced large improvements. As the system matured, the learning rate naturally declined. But it had not plateaued. Each month still brought incremental improvements: a new contract type the AI had not seen before, a regulatory update that required knowledge base revision, a client-specific preference that refined the output for that account.

The compounding effect was becoming visible in the firm’s economics. Average review time per contract had decreased from 47 minutes in the firm’s early months to 23 minutes. Error rates on the AI’s first-pass output had dropped from 8.7 percent to 3.2 percent. Client satisfaction scores, measured through quarterly surveys, had risen from 7.4 to 8.8 on a ten-point scale, driven largely by faster turnaround and more consistent quality.

These improvements flowed directly to margins. The firm's cost per engagement was declining while pricing remained stable, producing the margin expansion that the economic models predicted. From 22 percent EBITDA margins in the firm's earliest months to 32 percent now, not the 36 percent that Maya had described as the AI-native target, but moving in the right direction.

Sarah reviewed these numbers every week. She shared them with the team at the monthly all-hands meeting. She shared them with Alex at their biweekly check-ins. The numbers told a story of a firm that was learning, improving, compounding, doing exactly what the AI-native thesis said a firm should do.

But the numbers did not capture everything. They did not capture the confidence that Anika had developed, the way she could now scan a batch of AI outputs and identify systematic issues in minutes. They did not capture Nolan's growing judgment, his ability to spot the regulatory nuances that even the AI struggled with. They did not capture the conversations at the weekly retrospectives, where the team debated process improvements with an intensity that would have seemed bizarre at a traditional firm.

And they did not capture the absence of Joshua—the hole where his expertise had been, the complex regulatory matters that now took longer because Ryan and Megan had to cover his work alongside their own, the reminder that building a culture sometimes meant losing people you valued.



It was a Thursday evening in late October. Sarah had bought fourteen tickets to the Arizona Cardinals game against the Seattle Seahawks. Thursday Night Football. She did not particularly care for football, had never followed a team, and could not have named the Cardinals' starting quarterback if her life depended on it. But David was a devoted fan, Ryan had played in college, and several of the engineers had a fantasy league that consumed an unreasonable amount of their lunch breaks. She thought it would be good for the team to spend time together building bonds outside of just work.

She walked through the office at 4:30, handing out tickets and telling people to wrap up early. Most of them were already packing up. Megan grabbed hers

with a grin. Nolan high-fived Ryan. Karen and Wei, the two Legal Quality Analysts who had been with the firm since the pod launch, were already debating whether the Seahawks' secondary could handle the Cardinals' receivers.

Anika was still at her desk, headphones on, scrolling through a batch of AI outputs with the focused intensity of someone who had forgotten what time it was.

Sarah tapped on the glass. Anika pulled off her headphones.

“Game starts at 5:15. You’re coming.”

Anika glanced at her screen. “I found a pattern in the insurance disclosure language. I think we can flag it automatically. I just want to finish tagging the examples.”

Six months ago, Anika had been a recent law school graduate who could not get hired at a traditional firm. Now she was catching errors that experienced lawyers missed, proposing workflow improvements that saved the firm hours per week, and building judgment inside a system that most of the profession did not yet understand existed. Joshua would have said she was not yet a real lawyer. Joshua was gone. Anika was here.

“Don’t worry,” Sarah said. “It’ll be there tomorrow. Let’s go spend some time together. The bonds we build outside of work are what hold the pods together inside of it.”

Anika hesitated. The truth was she loved being at Candor. After eighteen months of rejection letters and interviews that went nowhere, she had found a place that valued what she could do rather than where she went to school or who she knew. Every morning she walked into the office grateful, and every evening she stayed late not because anyone asked her to but because she was terrified that if she eased up, even slightly, someone might realize they had made a mistake hiring her. The insurance disclosure pattern she had found was real, and tagging it would make the system better. Leaving it for tomorrow felt like a risk she was not sure she could afford.

Sarah read the hesitation. She had seen it before in people who had been underestimated. They worked harder than anyone because they believed they had less margin for error.

“Anika. Don’t worry. We love what you are doing and there is nothing that is going to fall apart at this job during this football game. I want you to go so you can strengthen the bonds with your colleagues. So please consider closing the laptop and joining us.”

Anika closed the laptop and grabbed her bag. “I don’t actually know anything about football.”

“Neither do I,” Sarah said. “David will explain everything whether we want him to or not. Meet everyone on the north side of the stadium. I’m buying drinks, food, foam fingers, t-shirts, whatever. The whole deal.”

She had bought the tickets because culture was not just about pod retrospectives and values statements on a whiteboard. Culture was about whether people wanted to be around each other when they did not have to be. Whether the engineer and the lawyer and the quality analyst could share nachos and argue about a third-down call and come back to the office on Friday morning with something that no org chart could produce. The McKinseys and Goldmans of the world understood this. Their offsites and team dinners were not perks. They were ‘social’ infrastructure.

Sarah was the last to leave the office. She looked back before locking the door. Every light was off. Every desk was empty. At 4:33 on a Thursday. She smiled. At her old firm, the associates would still be at their desks for another two to four hours, billing time, performing busyness, confusing presence with productivity. The culture she wanted was not one where people stayed late because they were afraid to leave. It was one where they worked hard, worked smart, and then went to a football game together.

They were building that culture. Slowly. Imperfectly. But they were building it.

---

COMING NEXT WEEK

Chapter 9  
*Measure Twice, Cut Once*

---

*Candor has made real progress but it is not always a linear path to building a lasting enterprise. Chapter 9 is about what happens when your AI-native system fails in front of the client who matters most. Eleven errors. A lost client. An \$80,000 rework bill. But the real story is what comes after: the rebuild. Candor OS 2.0 replaces rigid pipelines with AI agents that can reason, plan, and reach into regulatory databases through standardized interfaces. The client came back—not because Sarah pitched harder, but because the system she had now built was even more robust.*

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 Bucarius 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>