
The Great Decoupling

The Evolution of Services Firms when the Value of Code Is Infinite, but the Value of Coding Is Zero

April 2026

Frank D'Souza

Co-Founder and Managing Partner

Recognize

Table of Contents

INTRODUCTION	3
PART I: THE DECOUPLING.....	3
THE VALUE OF CODE AND CODING ARE PARTING WAYS	3
THE BURDEN OF ABUNDANCE: CHEAP CODE IS NOT THE SAME AS RELIABLE SOFTWARE	4
THE BOTTLENECKS HAVE MOVED	4
PART II: LESSONS FOR SERVICES FIRMS	6
NEW STRUCTURAL MODELS	6
NEW ROLES	7
THE THREE-PHASE TRANSITION.....	7
A CHALLENGING CULTURAL PIVOT: OVERCOMING HEADCOUNT-CENTRIC MUSCLE MEMORY	9
THE END OF PROGRAMMING, BUT NOT PROGRAMMERS	10
PART III: LESSONS FOR BUYERS OF SERVICES	11
CLOSING THOUGHTS	12

Introduction

In the 1990s, the advent of low-cost bandwidth and collaboration technologies ushered a generational shift into the technology services industry. The era of labor arbitrage allowed software developers in low-cost countries, such as India, to participate in the global services economy, dramatically changing the operating model across the sector.

Today, for only the second time in 40 years, the industry is on the cusp of another significant operating model disruption driven by advances in generative artificial intelligence (AI). By automating much of work historically done by humans, AI has become the "new India" for the industry. For services firms that traditionally generated revenue by providing human-intensive offerings, this raises a critical question: What happens when machines do the work that people once did?

In this paper, I try to answer that question. I focus on the subset of firms that develop code and software because today's large language models (LLMs) are extremely capable in these areas. But the lessons apply to the broader services industry. The same forces are at work.

I am writing for two audiences. The first is leaders of services firms who need a practical framework for navigating these shifts. The second is executives who buy technology services and are rethinking how to evaluate and procure providers.

As a longtime executive and now investor in technology services, I know the industry has weathered change before. But the AI services era will have impacts that prior shifts did not — on the profit pools, professions, and structure of the industry.

There is opportunity in this shift. Capturing it will mean abandoning many of the systems, incentives, and instincts that worked in years past.

Part I: The Decoupling

The value of code and coding are parting ways

To understand the impact of AI on the software services industry, it is useful to understand the diverging values of code and coding. The *task* of coding is becoming cheap, even while the output, well-functioning code, is worth more than ever. I call this the **paradox of value**.

Code is becoming one of the most valuable resources on Earth. As a primary tool for expressing human intent, code is the operating system of global finance, healthcare, infrastructure, and beyond. As software integrates further into our physical and cognitive worlds, the value of functional code compounds.

Coding — the manual labor of writing code — is moving in the opposite direction, shifting from a specialized craft to a commodity. LLMs can translate a business requirement into Python or Java at a marginal cost approaching zero, and AI can already automate virtually every stage of the software engineering lifecycle.

These lessons are reflected in the market, which is signaling loudly that software, especially AI software, will confer competitive advantage. In early 2026, OpenAI was valued at \$840 billion and Anthropic at \$380 billion — a feat that would have seemed inconceivable for software companies only a few years ago. At the same time, valuations of publicly listed technology services firms have declined over the past 12 months. Whatever your view of valuations, the markets are pricing in a world where leverage comes from code, data, models, and agents and not from talent-heavy service delivery.

For 40 years, we lived in an **era of coding scarcity**, where fluency in machine language was a natural bottleneck. We have entered the **era of coding abundance**.

Providers of software services must ground their thinking and analysis in this divergence. The output or outcomes they produce have never been more valuable, which presents a significant revenue growth opportunity. At the same time, the factor of production — that is, writing code — has never been cheaper, representing an opportunity for margin expansion.

The burden of abundance: Cheap code is not the same as reliable software

When something becomes cheap and abundant, the world reorganizes around it, as we have seen with storage, bandwidth, and compute. This reorganization creates new types of problems and new risks.

Code is no different; “good-enough” code generated by machines will soon be everywhere. But this creates a **paradox of complexity**: the easier it is to generate, the greater the surface area for risk. While the marginal cost of producing code approaches zero, the marginal cost of certainty — the guarantee that code is secure, ethical, fit for purpose, compliant, and integrated — rises.

In other words, abundant code is not the same as reliable software. Abundant code can create:

- **Technical debt:** AI can generate legacy code faster than humans can understand it.
- **An accountability vacuum:** If machine-generated code causes a systemic failure, who owns the recovery?
- **Security risk:** More lines of code mean more vulnerabilities to patch and monitor.

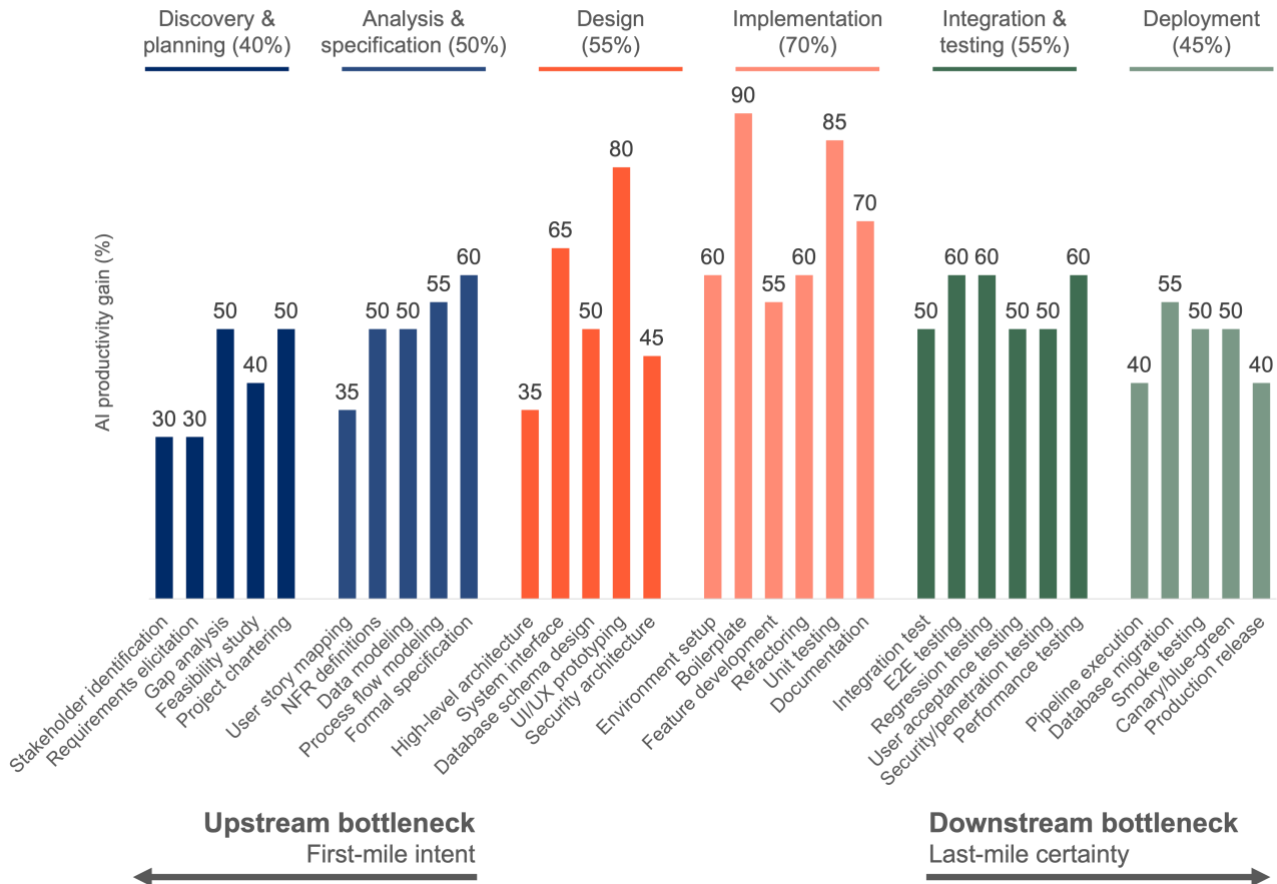
The bottlenecks have moved

In the era of coding scarcity, the bottleneck in software engineering was the manual labor of writing syntax, but AI has compressed that dramatically. The chart below shows AI’s productivity impact across the software engineering lifecycle. The gains are greatest in the “implementation”

phase, where boilerplate coding and unit testing have seen reductions of up to 90% and 85%, respectively.

AI's Uneven Impact on Software Development

Estimated time saved by task, across the software engineering lifecycle



But in the era of code abundance, the bottlenecks in software engineering have not disappeared. Rather, they have moved upstream and downstream.

Upstream, the bottleneck is now the **description of intent** — the conversion of messy, contradictory human problems into precise specifications that machines can execute. These are “first-mile” problems. Machines can build the house, but they don’t know what its future residents will want and need.

Downstream, the bottleneck is the **cost of certainty** — the effort to make code secure, ethical, fit for purpose, compliant, and integrated. This effort grows as the effort to produce code falls. These are “last-mile” problems. AI can generate legacy code faster than humans can understand it, creating pressure on long-term maintenance, accountability, and verification.

In a world where generating syntax is now nearly free, the burden of building reliable software systems remains high. Transforming machine output into reliable software requires effort because the bottlenecks have shifted from the task of writing code to the high-level challenges of defining "first-mile" intent and ensuring "last-mile" certainty. Ultimately, code may be abundant, but the human-led orchestration of that code into a trusted business outcome is the industry's new and most valuable frontier.

Part II: Lessons for Services Firms

As we enter the era of coding abundance, services firms must confront a fundamental truth: the market will no longer pay for the simple act of building, but it will pay more than ever for the impact of the build. This shift does not render the services partner obsolete; rather, it reinforces the classic theory of comparative advantage. Services firms will remain relevant and valuable when they can turn abundant code into trusted, secure outcomes faster and more reliably than a client can do alone.

New structural models

For services firms, several shifts are underway:

- The market will no longer pay for the work of building. Instead, it will pay for the **impact of the build**.
- The traditional talent model, rife with engineers, will give way to **hybrid teams** of highly specialized humans complemented by AI tools and agents.
- Firms that have **proprietary intellectual property (IP)** in the form of code, data, models, and agents will have an advantage.

To succeed, firms must abandon the old model and adapt to these trends with a new blueprint:

- **Effort-agnostic commercial models:** A firm will have trouble scaling if it bills by the hour while effort to deliver is falling by half. Firms must monetize outcomes, IP, or client usage. This lets a firm capture AI's efficiency gain, rather than passing it to the client.
- **Impact over headcount:** Firms must measure success not by bench strength or talent density, but by the ability of specialized hybrid teams to use AI to solve complex problems at speed.
- **Intelligent IP:** Generic AI tools are available to everyone. To differentiate, firms must build platforms that automate the routine work of the past and create moats of proprietary code, data, models, and agents.

New roles

These structural changes demand a different kind of workforce. In the scarcity era, the services firm provided the “hands” — people skilled in writing code and building the product. To survive in the abundance era, services firms must provide the “head” — people skilled at ensuring the end product delivers economic value. In practice, this means hybrid teams of humans and AI, with humans directing, reviewing, and validating what AI produces.

These new human roles cluster around the upstream, or “first-mile,” and downstream, or “last-mile,” bottlenecks described above.

Upstream roles:

- **The value orchestrator** ensures that what will be built adheres to a coherent business strategy and solves the right problems. In other words, humans must decide what to build and why.
- **The enterprise architect** ensures that what will be built fits within the existing systems of a large, complex enterprise — its security protocols, regulatory requirements, and legacy infrastructure.
- **The intent curator** turns business needs into specifications precise enough for AI to execute.

Downstream roles:



- **The results orchestrator** ensures that disparate AI-generated services integrate into a coherent, reliable business result.
- **The ethics steward** optimizes for equity, safety, guardrails, and long-term stability in an environment where machines may not have the context to make the right trade-offs.
- **The accountability steward** stands behind the output, bridging the “certainty gap.” When AI generates an error or a security flaw, the accountability steward owns both the fix and the outcome.

The three-phase transition

The main fear for services leaders is the J-curve — the period where AI-driven efficiency reduces billable hours and revenue before other commercial models can compensate. To navigate this, firms must adopt a prudent mix of three commercial models: pricing based on input, output, and finally outcomes.

New Commercial Models for Services Firms

Three pricing models to adapt in the AI era

	 Input pricing	 Output pricing	 Outcome-based pricing
What firm's client pays for	Time and materials	Fixed fee per deliverable	Share of business value
Success metric for firm	Margin expansion per billable hour	Growth in revenue per employee	Growth in long-term client lifetime value
Positioning example: health insurer	<i>"We provide the AI-native teams and expertise to build your claims system."</i>	<i>"We manage your claims operations for \$X per member per month."</i>	<i>"We take responsibility for reducing your medical loss ratio."</i>

Input pricing, with internal gains from AI

Firms continue to bill for time and materials but protect margins by using AI internally — for, say, boilerplate coding, unit testing, and documentation. The client sees a familiar pricing structure; the change is behind the scenes. Teams may be smaller, but higher hourly rates for the reduced team protects revenue and expands margins.

- **Objective:** Shift talent from routine work and toward high-value upstream and downstream tasks. Gain experience with AI tools and AI-native teams.
- **Metric of success:** Margin expansion per billable hour.
- **Example:** A services firm is hired by a health insurer to implement and maintain its claims platform. The engagement is billed by the hour, but internally, the firm uses AI for coding, migration, and testing, freeing key engineers to focus on integration and compliance. The services firm's positioning is, "We provide the AI-native teams and expertise to build your claims system."

Output pricing, with certainty bundles

Firms offer a flat fee for a defined output, with an assurance that the output is secure, fit for purpose, ethical, and compliant — what I call a "certainty bundle." The output metrics a firm chooses should align deeply with the volume drivers of the client's business.

The firm assembles a hybrid team of first- and third-party humans, and first- and third-party AI agents, who deliver the output efficiently and with reliability. If this team completes in 5 hours what used to take 40, the firm retains the value of the 35 hours saved.

- **Objective:** Decouple revenue from the clock, and use hybrid teams to deliver certainty.
- **Metric of success:** Growth in revenue per employee (RPE).
- **Example:** The services firm now offers insurers a certainty bundle: end-to-end claims management for a fixed price. The insurer buys a guaranteed operational output, rather than a team. The services firm's positioning: "We manage your end-to-end claims operations for \$X per member per month."

Outcome-based pricing partnerships

As stewards of outcomes, firms take a share of the business value they help create — through gain-sharing arrangements or usage-based pricing on proprietary tools and platforms. This model has been the holy grail of the services industry for decades. There are considerable frictions to adoption, which may be slow. However, as services become more IP-intensive, outcome-based partnership will make more economic sense for buyers and sellers. This is because the value the firm delivers with IP is harder for clients to replicate internally, and IP will make it easier to measure the firm's impact on business outcomes.

- **Objective:** Align the firm's incentives with the client's ROI.
- **Metric of success:** Growth in long-term lifetime value.
- **Example:** The same firm commits to running and improving overall operations. Compensation is tied to a measurable result, like the insurer's medical loss ratio. If processing and member experience become faster, more accurate, and less costly, both parties benefit. The services firm's positioning: "We take responsibility for reducing your medical loss ratio."

A challenging cultural pivot: overcoming headcount-centric muscle memory

A formidable barrier to this transition is the headcount-centric muscle memory of the industry's incumbent players. For decades, sales, delivery, and operations leaders have been incentivized to measure success by the scale and growth of their teams. This is reinforced by customers accustomed to auditing input hours rather than valuing the speed and integrity of an output.

Successful firms are those that decouple business growth from payroll growth. Sales leaders must pitch certainty instead of capacity; delivery leads must resist the urge to solve complexity by adding more people to a project. This transition requires firms to abandon the very metrics that defined success in the scarcity era, and replace them with structures built for the new era:

- **Incentive redesign:** Current compensation models inadvertently reward bloat by tying bonuses to team size or contract volume. Sales commissions should move away from total contract value based on headcount and toward **revenue per outcome** delivered.
- **Leadership signaling by ‘burning the boats’:** This transition will not happen bottom-up, and will instead require a visible, top-down rejection of scarcity-era metrics. For instance, a firm might formally retire utilization as a primary success metric and replace it with **outcome density** — the business value delivered per team member per quarter. This signal forces the organization to prioritize certainty over capacity.
- **Client re-education:** Buyers are often complicit in the old model, using timesheets as their only known control mechanism. Services firms must proactively lead their clients toward new evaluation frameworks — for example, via **certainty dashboards** that visualize progress on business milestones, security, and reliability, rather than human hours expended.

The end of programming, but not programmers

An evolving discipline: from syntax to systems

In the scarcity era, "programmer" and "software engineer" were used interchangeably. The bottleneck of engineering was the manual, syntax-heavy labor of writing lines of text, and the main source of failures were bugs caused by human error. As a result, formal approaches (e.g., SDLC, CMM, Agile) focused on making software engineering repeatable and fast, while ensuring that the end result conformed to user requirements.

In the abundance era, software engineering will become a discipline focused on outcomes. The hybrid engineering team will be skilled in expressing intent, architecting and steering, reviewing and integrating, and building trustworthy systems in a world of machine-generated code.

Hybrid teams

Engineering teams will evolve from human-only Agile squads to **hybrid teams**, comprised of upstream and downstream roles, like enterprise architects and accountability stewards, who work alongside AI agents.

In the scarcity era, a services firm might have staffed a 45-person team for 18 months to migrate an insurer's claims platform. In the abundance era, an orchestrator leads a hybrid team of six people and a suite of AI agents. In any week, they might work to understand how technology can provide the customer with competitive advantage, mediate between the client's compliance team and the AI's recommendations, ensure that the system being built will integrate safely with other client systems, and run an audit to ensure outputs are explainable to regulators.

The death of the pyramid

In the traditional services model, the talent structure was shaped like a **pyramid**. At the base, massive cohorts of entry-level engineers were hired to do manual coding and other lower-level tasks. But AI makes this structure obsolete. Firms can no longer scale by billing hours for hands to build products, and the J-curve transition requires a shift away from labor-intensive delivery.

In the AI-native world, a team of humans will resemble a **diamond** instead of a pyramid. In the broad middle of the diamond, enterprise architects, intent curators, value/results orchestrators and ethics/accountability stewards manage AI output. At the narrow base, a smaller, highly curated group of apprentices learn to steer AI agents and fill the accountability vacuum.

A shrinking base means fewer entry-level jobs in an industry that has employed millions. The challenge, then, is also one of workforce development: How do firms cultivate the next generation of senior talent when there are fewer junior roles to serve as the traditional training ground?

The answer is strong apprenticeship — smaller cohorts trained from the start to fill upstream and downstream roles. The industry as a whole would benefit from treating this as a shared responsibility, including by finding commercial models that can sustain the ongoing cost of apprenticeship.

Part III: Lessons for Buyers of Services

The “Great Decoupling” of code and coding alters the relationship between enterprises and their services partners. The procurement strategies and evaluation metrics that CXOs and Boards have used for 40 years — rate cards, headcount audits, utilization tracking — are becoming liabilities. Technology buyers must move beyond labor arbitrage and hourly billing to output or outcome stewardship.

Output or outcome specifications, not staffing requests

In the scarcity era, buyers requested capacity — e.g., "30 Java developers for 12 months." In the abundance era, buyers should specify business results — e.g., "Reduce claims processing time by 40% with a full security and compliance guarantee." The firms that win will be those that can commit to certainty bundles and stand behind the result, regardless of the effort.

The competition will be fierce, as firms compete not only with one another, but with clients' own internal teams, i.e., Global Capability Centers. If an internal team can use ubiquitous AI to generate “good-enough” code, the bar for external firms rises.

The certainty scorecard

Procurement teams need a new evaluation framework to replace the traditional rate card. Instead of evaluating bench strength and hourly rates, buyers should evaluate services firms based on a structured "certainty scorecard, which might measure four dimensions:

- **Outcome reliability:** A track record of hitting business milestones, not just deploying code.
- **Proprietary IP:** The depth of the firm's own code, data, models, and agents — tools that go beyond commodity AI.
- **Security posture:** The ability to manage the security risks inherent in machine-generated code at scale.
- **Accountability:** Willingness to be an accountability steward when AI generates an error or systemic failure, both correcting errors and owning the outcome.

Closing Thoughts

Unsurprisingly, many leaders in the technology services industry are watching the decoupling of code and coding with great trepidation. I understand this. After firms spent decades perfecting the structures and systems of the scarcity era, the abundance era can seem daunting.

But in my own career, I have watched the industry absorb many waves of change — offshoring, internet, and cloud, to name a few. Each time, the firms that adapted early found more opportunity on the other side than they expected.

For now, AI is another tool — a powerful one, certainly, and one that will change many industries more than any single innovation before it. But it is a tool nonetheless. The value of a hammer dropped when the nail gun was invented, but houses are more valuable than ever. The same is true of code. Tools change as industries grow, and new tools serve the people who learn to use them well.

For services firms, that means building for certainty, staffing hybrid teams, and measuring success by output and outcomes. For buyers of services, it means demanding more from their partners and evaluating them differently. Those that do will find partners worth paying for.

And for both services firms and buyers, the question for leaders is no longer, "How do I produce more code?" but "How do I turn abundant code into trusted outcomes faster than anyone else, without losing integrity or security?"

The opportunity is substantial, but the metrics and instincts of the scarcity era will not carry the industry through the abundance era. Success belongs to those willing to change.

DISCLOSURES

Frank D'Souza is a Co-Founder and Managing Partner of Recognize Partners LP ("Recognize" or the "Firm"), a private equity firm investing exclusively in digital services companies. This opinion piece (the "Information") reflects his own views as of the date hereof and does not necessarily reflect the views of any entity he represents; neither Mr. D'Souza nor Recognize undertake to advise you of any future changes in the views expressed herein. The Information is being furnished to the recipient ("you") solely for informational purposes. This Information is not, and may not be relied on, in any manner, as legal, tax, investment, accounting or other advice. This Information is not, and should not be construed as, an offer of investment advisory services by Recognize, nor as an offer to sell (or a solicitation of an offer to buy) an interest in any investment sponsored by or affiliated with Recognize. The Information is not an endorsement of any particular security or investment opportunity.

Certain information contained in this Information has been obtained from published and non-published sources prepared by other parties, which in certain cases has not been updated through the date hereof. While such information is believed to be reliable for the purposes of this Information, neither Mr. D'Souza, nor Recognize, nor their affiliates assume any responsibility for the accuracy or completeness of such information, and such information has not been independently verified. Statements contained in this Information (including those relating to current and future market conditions and trends in respect thereof) that are not historical facts are based on current expectations, estimates, projections, opinions and/or beliefs. Such statements involve known and unknown risks, uncertainties and other factors, and undue reliance should not be placed thereon. In addition, no representation or warranty is made with respect to the reasonableness of any estimates, forecasts, illustrations, prospects or returns, which should be regarded as illustrative only, or that any profits will be realized. You should make your own investigations and evaluations of the contents of the Information and should note that such information may change materially.