



Software Unbound

The Agentic AI Reset




www.7t.ai

Executive Summary

Since the rise of Agentic AI rarely does a day go by when people don't talk about its immense potential. However, lately this excitement has changed into doomsday predictions of massive job losses, with statements like "Software is Dead" and "SaaSocalypse is Coming". While these declarations are provocative, they can't just be dismissed outright. They do deserve closer examination to understand the nature of this impact — to separate reality from paranoia. Let us face it, the software industry is fundamentally re-architecting itself — from the way code is created all the way to how it is deployed, licensed, and sustained. This change is at the root of this disruption. This article takes a deeper look at industry dynamics and explains the true nature of this disruption. We examine questions such as: How will the change happen? What is so different about Agentic AI and what will be its impact? How will jobs be disrupted? Who is at the most risk of displacement? We address these questions and make the argument that 'software' will continue to flourish but it won't be without a new set of winners and losers. We are moving from an era of scarcity (limited by software engineering powered by human resources) to an era of abundance (limited only by human intent). The basis for this outlook includes history of technological innovations in computer hardware and software, the latent demand for digital solutions that could not be fulfilled without AI, and AI as the force multiplier of human ingenuity.



Vibe Coding isn't about 'one and done'

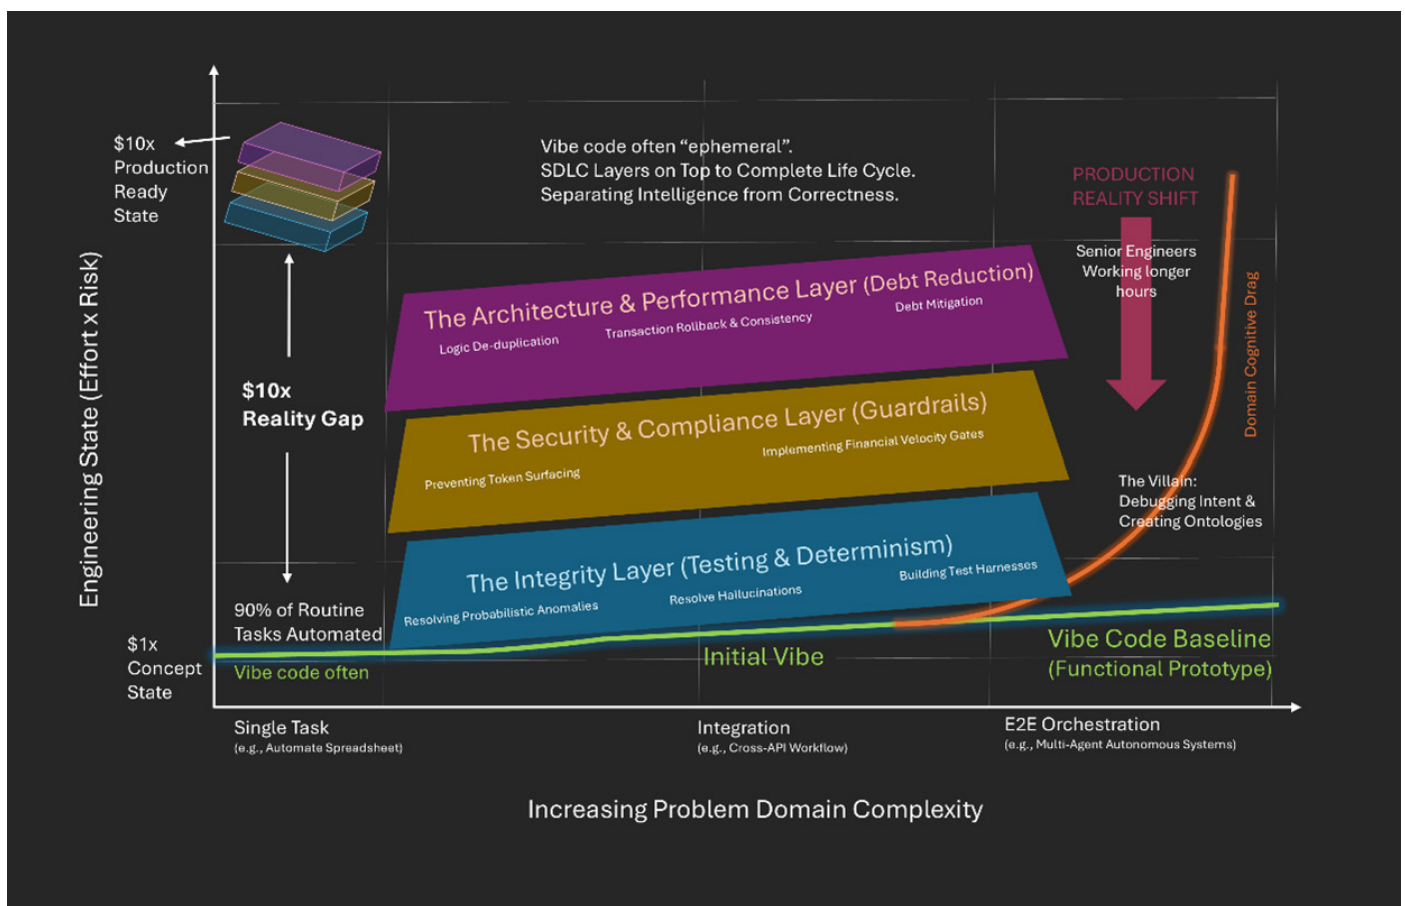


Every aspect of how software is created or modified, tested and validated, deployed into production, and supported is changing. In some cases, large portions of this value chain have already changed due to "vibe coding", a term coined by Andrej Karpathy used to describe how developers use Agentic AI built into popular IDEs (Integrated Development Environments) to generate code through well-defined prompts that describe requirements, constraints, and contextual information. The better software engineers do in describing these elements, the better the output. The quality of code generated by Agentic AI within minutes is better and with fewer bugs, than what the above-average programmer would produce in a significantly longer time (5x-10x). However, it is important that the requirements, constraints, and contextual information specified in the prompts are written by a seasoned software engineer. It is possible that in the future this requirement may not be as stringent as it is currently, but that depends upon future innovations in coding agents. As of the time of this article and based on first-hand experience with vibe coding leveraging Claude Opus 4.5 and Cursor IDE, a senior software engineer still needs to review all generated code to make sure it does what was intended. Most often an imperfect or ambiguous set of prompts is the reason why generated code does not meet the mark, but it could also be due to limitations in Agentic AI training. It is hard to distinguish between the two, but in our experience, it is possible to get AI agents to generate correct code after multiple iterations of prompts refining our intent each time. In summary, the vision or idea and how it is expressed in requirements, followed by their expression in software terms (the architecture, constraints, environment) determines the quality of output produced by vibe coding. We aren't at "auto pilot" yet. We are at "augmented pilot."

While vibe coding can be used in many stages of the SDLC (Software Development Life Cycle), the initial benefits are what has caught the industry's attention with most arguments about an impending "SaaSocalypse" and "Software is Dead" based on just this aspect. If you can build it 10x faster and at 1/10th the cost, why buy expensive software licenses and subscriptions? Just build it yourself — driving small, medium, and large enterprises to stop paying for licenses and subscriptions. Well, not so fast! Creating the initial baseline code isn't all there is to the SDLC, the harder part is what happens after the initial version is created and how complex the problem domain is. In fact, reports are emerging in the industry [REF-1] that senior software engineers leveraging Agentic AI and vibe coding are working longer hours with fewer resources. Wait! That sounds like a contradiction — well it certainly is contrary to the benefits that Agentic AI and vibe coding were supposed to provide. Let's examine why and understand how the economics of the software industry are changing.

Figure 1 illustrates how Vibe Coding only provides the initial baseline and what it takes to create production ready software. The "Vibe Code Baseline" (the bright green curve) represents the generation of functional, prototype logic, which Agentic AI can now achieve with near-zero friction. But this baseline sits below the surface; to complete the software development life cycle (SDLC), engineering teams must add consecutive layers of complexity and validation, all while struggling against the rising orange curve of Cognitive Drag (the complexity of the domain). Concept State is the moment the agent outputs the first runnable prototype. This is "vibe coding." Production-Ready State is the state required to ship software. Achieving this is where the "longer hours" come from.

FIGURE-1: The Vibe coding reality gap



The SDLC Layers on Top

The total effort required to complete the life cycle (the area above the green line) is substantial, especially in high-complexity domains:

- **Layer 1:** The Integrity Layer (Testing & Determinism): Agents are probabilistic, so you need to prove the output always works. This requires more effort than the vibe coding itself to build deterministic test harnesses, integrate CI/CD, and resolve hallucinations.
- **Layer 2:** The Security & Compliance Layer (Guardrails): You don't just let a non-deterministic agent write the payment logic. This layer requires strict data privacy controls (preventing "token surfacing"), and recursive prompt security audits.
- **Layer 3:** The Architecture & Performance Layer (Debt Reduction): Vibe code is often "disposable." If this needs to live for more than 48 hours, a senior architect must perform logic de-duplication and remove any technical debt to prevent the "technical debt" drag that occurs when complexity spirals.

Why Senior Engineers Are Working Longer Hours

The vertical orange curve—Domain Cognitive Drag—is the real villain of this forecast. Vibe coding makes solving a simple problem like "Automate a Spreadsheet" effortless. But as the "Increasing Problem Domain Complexity" rises (moving right on the X-axis), the problem stops being about syntax (e.g., writing Java) and starts being more about End-to-End Orchestration. A senior engineer in 2030 won't be writing code; they will be creating ontologies that define knowledge and context for agents to act on and debug intent. The code generation is instant, but creating knowledge systems, context, and system validation in complex environments is profoundly manual and difficult. In fact, in a recent interview, AWS CEO Matt Garman stated, "Writing Java code snippets on demand will matter less, building end-to-end applications, understanding what customers actually need, and working fluently with cloud services will matter more."

Reallocation of Spend

Historically, lowering the cost of a resource increases its consumption (Jevons Paradox). Remember open-source software (OSS) a little over two decades ago? At the time some in the industry were all in and felt that was the future of how software would be built. While others were mostly on the fence and continued to build proprietary software. In the long run what we have now is software that is built faster, some fully open source and some proprietary, although a lot of the non-open-source software leverages open-source software components or tools. Also, Everest Group analysis from real world examples suggest productivity in SDLC has improved by 12% on average just in 2026 [REF-8]. When you combine both the stories, the result is that we will have more software that is built at a lower cost. Agentic AI and vibe coding will have an enormous impact on the industry as they drastically lower the cost of software.

There will be more software solutions due to gains in productivity than ever seen before. In addition, there has always been a latent demand for more software solutions that could never be fulfilled by industry without Agentic AI. For the first time there is a promising technology that will drastically change the economics of software by changing the DNA of the creation process. Software is poised to grow, not shrink.

The Evolution of the Cost Equation

The cost of running software has come down significantly over the last two decades through successive waves of technological innovation in hardware, operating systems, IDEs, reusable architecture patterns, APIs, test automations frameworks, and the availability of open-source components. But up until the arrival of Agentic AI and the ability to state requirements in natural language, we had pretty much hit a cost wall (see **Figure 2**) as the industry had run out of levers to reduce overall cost. Software still had to be written and maintained by humans (software engineers). Now the very DNA of software creation and maintenance is changing due to vibe coding, and this is driving down the cost of software again. Let's examine the cost of application software and how it has evolved.

Initial Cost Equation

Total Cost = Physical Hardware Cost + Human Cost (Owner + PM + BA + Arch. + Design + Dev + Test + DevOps + SRE) + Software License/Subscription Cost

Cost reduction through hardware virtualization & Test Automation

Total Cost = Fractional Physical Hardware Cost + Virtual Machine Cost + Human Cost (Owner + PM + BA + Arch. + Design + Dev + Fractional Test + DevOps + SRE) + Automated Test Cost + Software License/Subscription Cost

Cost reduction through OS virtualization

Total Cost = Fractional Cost of Physical Hardware + Fractional Cost of Virtual Machine + Cost of Container(s) + Human Cost (Owner + PM + BA + Arch. + Design + Dev + Fractional Test + DevOps + SRE) + Automated Test Cost + Software License/Subscription Cost

Cost reduction through DevOps/SRE Automation

Total Cost = Fractional Cost of Physical Hardware + Fractional Cost of Virtual Machine + Cost of Container(s) + Human Cost (Owner + PM + BA + Arch. + Design + Dev + Fractional Test + Fractional DevOps + Fractional SRE) + Automated Test Cost + Automated SRE Cost + Automated DevOps Cost + Software License/Subscription Cost

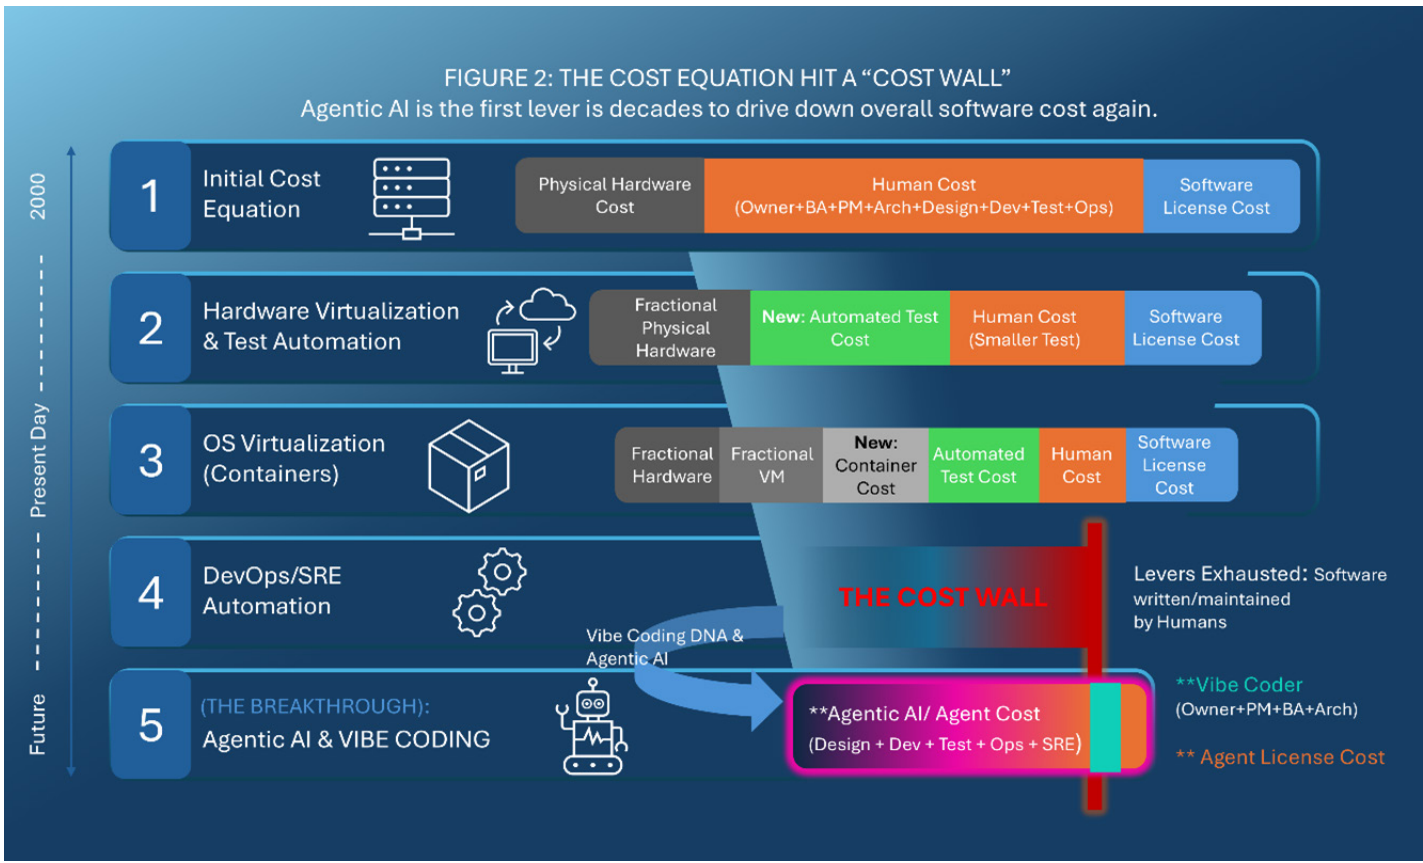
Cost reduction through Software Development Automation – Vibe Coding & Agentic AI

Total Cost = Fractional Cost of Physical Hardware + Fractional Cost of Virtual Machine + Cost of Container(s) + Human Cost (Owner + PM + BA + Arch. + Vibe Coder + Fractional SRE + Fractional DevOps) + Agent Cost (Design + Dev + Test + DevOps + SRE) + Software (User) License/Subscription Cost + Agent License/Subscription/Token Cost

The last equation introduces costs associated with Agentic AI. The drop in Software (User) License/Subscription cost and human cost will be compensated by Agent License/Subscription/token costs. We will see a shift where companies pay for **Digital Labor** instead of just **Seats**. The cost of Software, software license or subscription costs that are charged by product companies will go down as the cost of creating and maintaining software goes down. However, the number of licenses or subscriptions required would go up due to a new crop of users — Agents. CNBC producer Jasmine Wu coined a phrase that captures this shift nicely — SaaS (Software for agents as a Service).

In addition, there is another question of whether PM and architecture costs will become agent costs. Maybe, but it may take a few years to achieve that level of maturity. Everest Group research suggests a shift from labor-based roles to a combination of agent-driven execution, reusable asset leverage, and human-led orchestration. [REF-7] While there are stories in the press about AI startups building everything and operating everything purely using AI agents, those claims need validation given how nascent they are at this point. Such approaches need to provide answers to questions such as: When things go wrong, how will the founder solve them? When agents hit a roadblock in their ability to solve a problem, who will be there to break the logjam? Can other Agents help solve these problems? It explains why senior engineers are currently working longer hours than ever before due to the need to inspect/correct code created by agents or to ensure agents produce consistent outputs at run time. Let's now look at exactly what the capabilities of AI agents are.

FIGURE-2: The cost equation hit a 'cost wall'



AI Agents Are Intelligent Software Robots

Prior to the arrival of agents to the AI scene, we had LLMs. You could do a lot with LLMs through prompt engineering, but ultimately, they provided answers to user queries quickly and in many cases surprisingly well. This came as a surprise not just to the users, but even to the creators of the models. However, they did have the tendency to make things up (hallucinate), although (in reality) LLMs were simply finding the best statistical match while responding to queries based on their past training (on mountains of data) that was optimized through RLHF (Reinforced Learning through Human Feedback).

Agents introduced a new paradigm emulating how humans solve problems through reasoning and executing actions according to a plan. Instead of simply responding to a user query based on the best statistical match and RLHF training, the idea of generating a plan to tackle the user's request and then following the best path (that could optionally use a known set of tools) before responding with an answer produced better results. This heralded the beginning of Agentic AI. Reasoning (building a plan using the LLMs ability to generate text) coupled with a well-defined role and an ability to take actions using available tools until the desired objective was met is how agents operated independently to accomplish tasks. To perform complicated tasks with many different sub-tasks a "crew" of agents could be used, with each member of the crew performing a sub-task, while another agent was responsible for evaluating the results to determine if the task had been performed satisfactorily to meet the overall goal, or whether more needed to be done. Agents have been quite successful in performing tasks where it is possible to clearly and easily determine whether a task succeeded or failed. This is the reason why coding agents have become so popular for programming tasks — the program either works or fails.

Now, a new crop of Agent-Native Foundation models (ANFMs), such as the ones pioneered by Claude Opus 4.6, the capabilities of LLMs have been optimized to provide agentic functionality natively. These ANFMs have improved fragmented software engineering workflows by shifting from a reactive chat assistant to a persistent, goal-oriented engineer. The integration of **Adaptive Thinking** (model's ability to adjust its logic and processing speed based on the complexity of the task) and **Extended Thinking Mode** (a high-compute state for the model to "ruminate" on complex problems before answering) allows the model to selectively "pause and plan" for up to 8 hours on architectural problems without over-taxing tokens on trivial syntax, while **Context Compaction** (model's ability to shrink massive datasets into core insights without losing essential nuance) and the **1M-token window** (which allows the model to process and recall information from thousands of pages of code at once) prevent "context rot," ensuring it retains a perfect mental map of a code repo over multi-day sessions. Crucially, by providing **Native Primitives** (which are built-in tools to enable functions like code execution and file handling that are "baked in" rather than added on) with secure write-access to repositories and terminal environments, ANFMs move beyond "code snippets" to execute end-to-end business cycles—autonomously diagnosing failures, refactoring across files, and validating its own PRs—essentially closing the loop between a high-level business requirement and a production-ready deployment. While ANFMs hold a lot of promise, at this point it remains unclear how much these ANFMs will reduce the burden on senior engineers.

The Path Forward

There is no doubt that the software industry is going through a radical shift, but it's an evolution, not an extinction. As the cost of creation drops, we won't see less software, we will see an explosion of it. This is a massive productivity play that will feed the latent demand that has always been there. We illustrate these ideas in **3.1**, **Figure 3.2** and **Figure 4**.

Figure 3.1 shows that in the Agentic AI age we will be able to innovate more and solve more complex problems with fewer engineers compared to the Manual age of the past. Also, low complexity problems will require significantly fewer software engineers in the Agentic AI age compared to the “cost wall” of the Manual age. As we look into the future (see **Figure 3.2**) the composition of all software authors will change to include citizen developers, product owners and product managers. This is what we call the democratization of software and innovation. As of 2026 we are already at a 1:1 ratio for professional software engineers and deployed agents. Deployed agents are expected to grow exponentially to increase to over 1 billion agents by 2029. When you combine citizen developers, software engineers, agents, and product owners/managers you have a very large development community resulting in more software solutions than ever before.

FIGURE-3.1: The Explosive Growth of Innovation

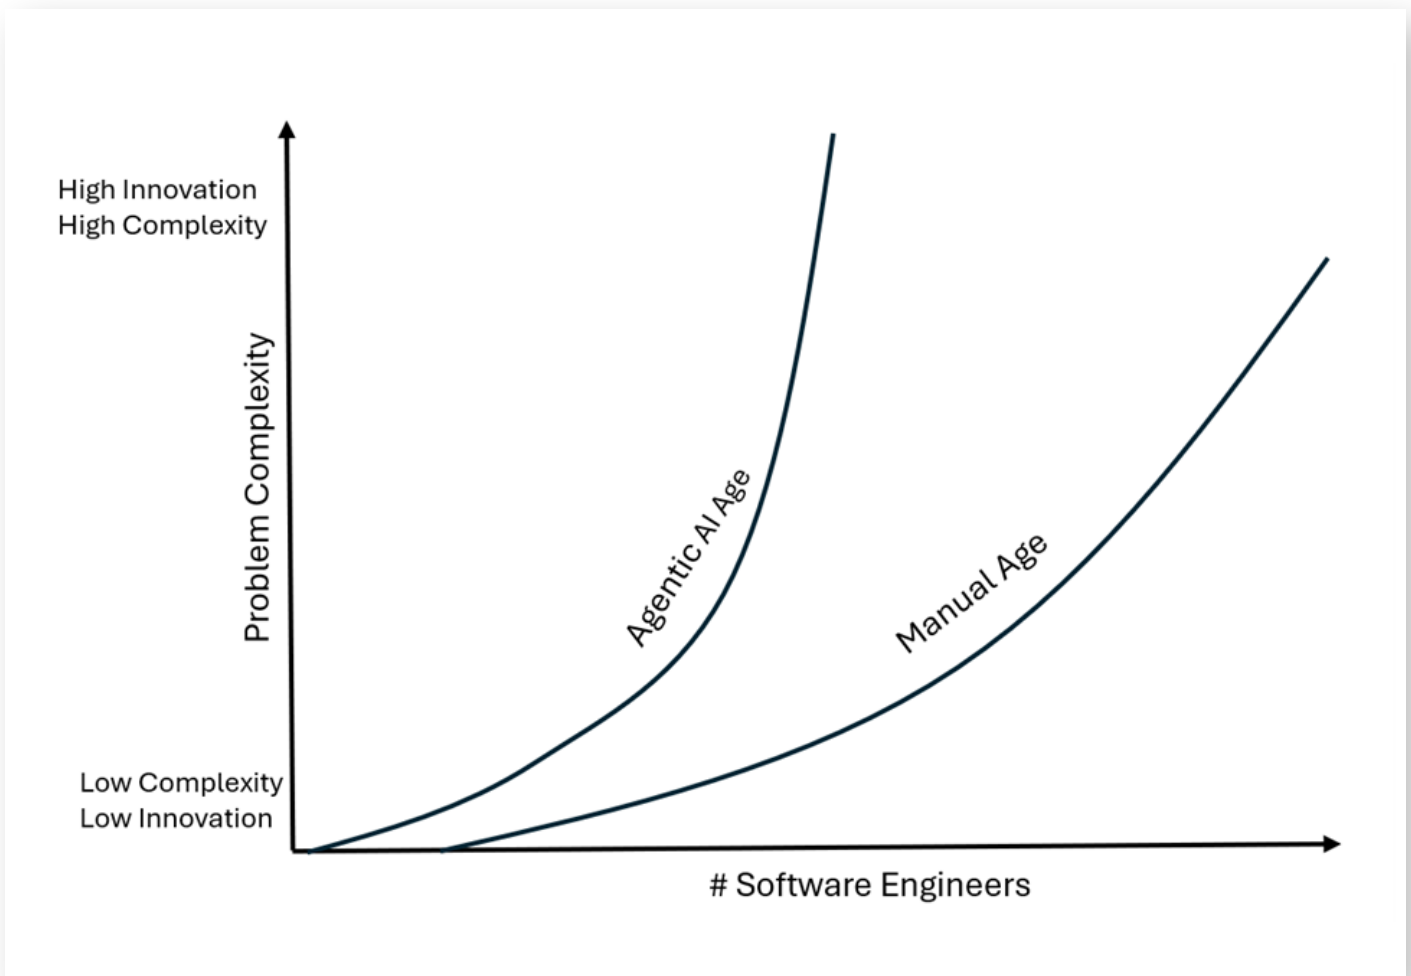


FIGURE-3.2: The Democratization of Software

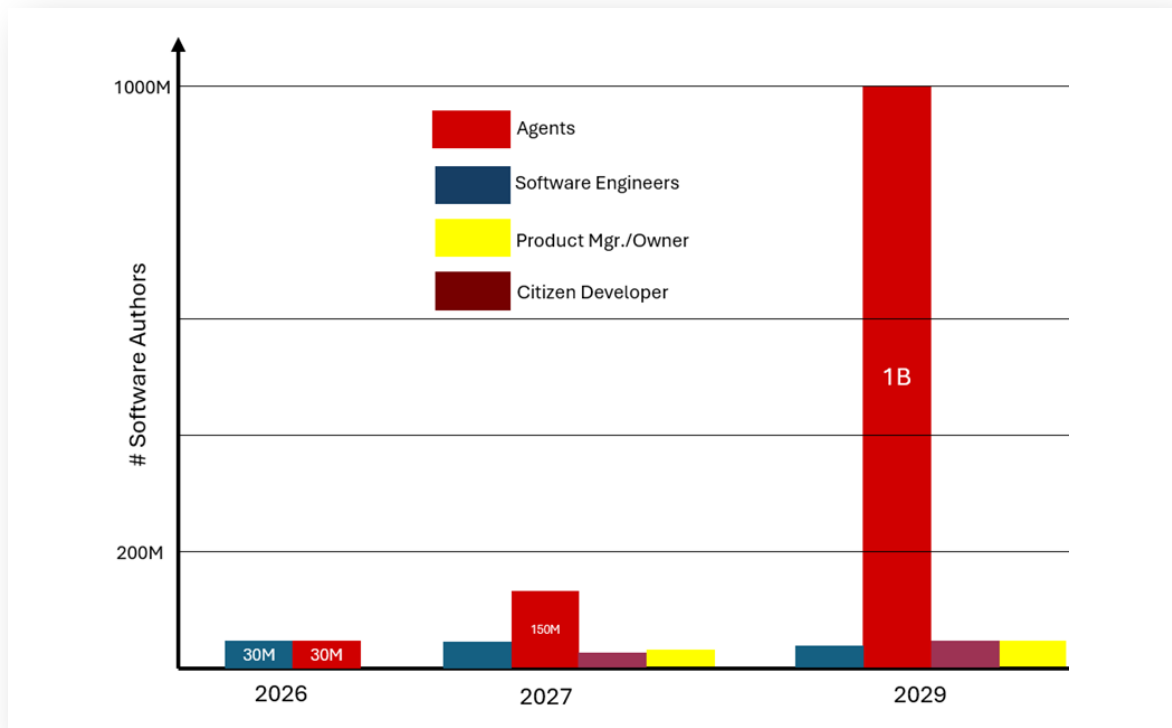


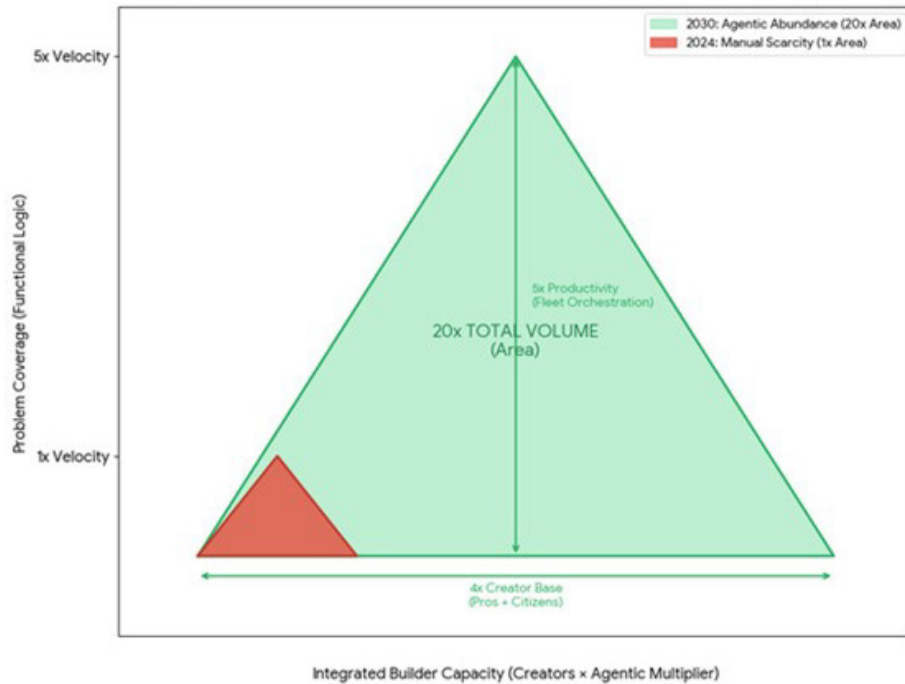
Figure 4 supports our projection of the upcoming software abundance with some quantification. The red triangle represents the **2024 Baseline (The Manual Era)**, where software production is limited by a small pool of professional developers and the speed of manual development. The massive green triangle represents the **2030 Target (The Agentic Era)**. Let us now list key contributing factors to this structural shift in who builds and how fast they build. The math isn't about "better tools"—it's about a 4x increase in the builder population combined with a 5x increase in velocity.

Key Contributing Factors:

- **From Scarcity to Abundance:** Up until 2024, software was an expensive, artisanal craft. By 2030, it will be a high-volume, automated utility.
- **The Role Pivot:** The 4x growth in "Builders" is driven by the collapse of technical barriers, shifting development from "Syntax-first" to "Logic-first"—effectively rendering the attrition of traditional "code-jockeys" a minor, inevitable rounding error in the face of massive Agentic AI scaling.
- **The Force Multiplier:** The 5x height growth represents the transition from a "human writing functions" to a "human overseeing autonomous agent cycles" that handle code, testing, and deployment.

Market analysts, AI vendors, and Strategy Consultants all share similar outlooks in how they see AI's role in the expansion of software. According to Gartner and IDC [REF-3, REF-5], we are moving from the world of 'Syntax Scarcity' to 'Agentic Abundance.' When the cost of building software drops by 95% (per Anthropic's productivity benchmarks), the volume of software produced will naturally expand by 20x to fill the vacuum of unsolved business problems. McKinsey & Company in their 2026 "Economic Potential of Generative AI" [REF-4] update estimates that agentic AI could add \$2.6 trillion to \$4.4 trillion annually to the global economy. This level of value creation is mathematically impossible without a 10x–20x explosion in the volume of functional logic (software) being deployed.

FIGURE-4: The Software Abundance Shift 2024 vs 2030



Admittedly, all this creates a "cold start" problem for new graduates. When AI functions as a pre-trained, lightning-fast intern, the entry-level bar is raised significantly. However, for those who can move beyond "vibe coding" into system architecture and AI orchestration, the opportunity has never been greater.

We think the cloud companies will make the most out of this shift. Today they are all a combination of IaaS, PaaS, and SaaS and have mature AI portfolios. As more businesses adopt agentic AI, they will want to implement some of the simpler capabilities in-house. For this they will leverage the AI models, hosting and management capabilities offered by the cloud vendors. On the other hand, pure-play software companies offering simple tools will face some initial headwinds and will need to come up with creative licensing models or verticalize their offers (vertical SaaS) to be industry specific to account for the shift in the user base from fewer humans to more agents.

Jensen Huang recently said that AI will largely leverage existing, market-proven tools rather than reinventing them [REF-2]. Andrew Ng in one of his newsletters — "The Batch" also feels that the Jevon's Paradox will apply to Software. One could argue that both industry leaders have too much vested in the adoption of AI, but when one reflects on these statements a little bit deeply, the conclusion seems obvious — Software is poised for growth!

Why some SaaS will survive the "Build vs. Buy" Trap, but not all of it

The argument that "anyone can now build their own ERP with AI" ignores three structural realities:

- 1. The Complexity Barrier:** You cannot "prompt" your way into a global Telecom billing system. These systems require architectural foresight that AI—which functions on statistical probability, not sentient ingenuity—cannot yet replicate. The "cost" of building is far from zero; you still require expert humans to audit generated code, design robust test data, and manage integration. Today's AI capabilities still demand a substantial—and expensive—human-in-the-loop.

- 2. Focus vs. Ownership:** Most enterprises want to run a business, not a software house. Even if building is "free," **maintaining** it is expensive. SaaS is a way of **outsourcing liability and compliance**. Even if the cost of development drops, the opportunity cost of managing a custom-built tech stack remains high. While niche, non-existent tools might be built in-house via AI, businesses will continue to subscribe to established SaaS for complex systems to avoid the burden of development
- 3. The Burden of Maintenance:** Software is a living organism. It needs security patching, API updates, and scaling. Most companies will still pay a premium for a "black box" solution that just works rather than managing a fleet of custom-built AI bots.

Enterprises also have to move slowly and carefully. They know that their mission-critical systems are part of too complex a web of many integrations, to be rewired hastily. So, despite the many advantages of Agentic AI it will take enterprises 12-18 months to gain from its benefits. So, the software industry is in no imminent danger of collapse.

The Winners and Losers

| Segment | Impact | Strategy |
|---------------------|------------------|---|
| Cloud Providers | High Growth | They win by hosting the models and the increased compute demand. |
| Product (SaaS) Cos | Stable/Evolution | Must pivot from "per-user" pricing to "per-agent" or "value-based" pricing. |
| System Integrators | High Disruption | Demand for "bodies" will crater. They must pivot to building platform-driven models to remain relevant. |
| Staffing Agencies | Severe Risk | The need for junior/mid-level "coders" will vanish. Revenue will likely shrink as teams become smaller and more elite. |
| Consulting Services | Some Disruption | Will need to demonstrate a greater sense of responsibility in advising clients amidst the increased set of options and opportunities. |

Impact on ISVs (SaaS or non-SaaS): They will not see a reduction in demand as hardly anybody wants to reinvent the wheel nor would they want to take on additional responsibilities to maintain, upgrade, and secure the software deployments. The economies of scale achieved by product are unchanged. Concerns about loss of revenue due to reduced user subscriptions are premature. ISVs will adjust their licensing models to account for agents in place of users. In addition, complex SaaS software, for example, billing software is rarely licensed on user subscriptions alone. Quite often it is a combination of multiple factors including, base platform fees, percentage of revenue billed, and other criteria. ISVs will also improve margins by lowering development costs using Agentic AI and limiting human costs to Owner, BA, and Vibe Coders.

Impact on Services Companies (Systems Integrators): There will be a major impact on the demand for their services. Firms that continue to rely on scale and labor arbitrage risk declining relevance, while those that invest in reusable assets, orchestration capabilities, and AI-native delivery models will define the next phase of the industry [REF-6]

Impact on Consulting Services Companies: The high-end consulting services companies will need to navigate carefully. Existing approaches to advisory services will need to evolve by providing more proof points, while at the same time augmenting Agentic AI based research with additional evidence based on recent company experience and tribal knowledge.

Impact on Staffing Companies: As the cost of building software drops significantly, requiring only Owner, PM, BA, and Vibe Coders, staffing software companies will experience significant disruption. Staffing companies may have the opportunity to provide Vibe Coders, but the loss of many junior and mid-level resources will lead to a significant reduction in revenue and margins. There may even be cost pressure on them to provide Vibe coders at lower rates, impacting US revenue. Software staffing companies are likely to be the hardest hit segment of the industry.

Our advice to staffing companies is to build an AI strategy to adapt to this change and start implementing it as soon as possible. We recommend the following three-pronged approach:

- Shift focus to senior specialist resources — architects and senior developers with critical thinking and strong communication skills. These are the resources that will be in demand as vibe coders and reviewers of AI generated code. Instead of seeking generalists, focus on hiring specialists in industry verticals that are in your sweet spot.
- Immediately start the process of removing any debt from the books. Debt can become the biggest issue when free cash flow starts to fall. The shift from higher margin lower cost resources to lower margin specialist resources will be gradual. It is more likely that lower cost higher margin resources will drop faster than the uptick in senior specialist resources as enterprises adapt to this age of agentic AI.
- Rebrand and Focus on Software Services: Invest in new subsidiaries that focus on higher-value software services. Develop re-usable frameworks leveraging agentic AI and then approach existing customers, offering software services instead of staff augmentation.

Conclusion: Productivity and Boom in Software, Not Doom

Software isn't dying. It's just going through a phase of disruption born from the Agentic AI innovation. This will translate into more software for at least two distinct reasons, if not more. One is to create software that is nearly impossible to build without Agentic AI, and the other is software built effortlessly by non-programmers. Take the case of O-RAN (Open Radio Access Network) where the complexity of the network is so high that troubleshooting network problems manually is impossible. Using Agentic AI, it is possible to build “intent-based” network management software, that optimizes and proactively troubleshoots issues based on “intents”. Non-programmers can now build custom tools like a travel blog that can be updated from anywhere in the world.

As the software industry rearchitects itself, some businesses caught on the wrong side of this change and solely reliant on labor arbitrage will suffer the most. ISVs (including the Agentic AI enabled new crop of startups) and

cloud companies will see tremendous opportunities for growth and innovation. There is also a "cold start" problem for new graduates. When AI functions as a pre-trained, lightning-fast intern, the entry-level bar is raised significantly. However, for those who move from "writing code" to "orchestrating systems," the opportunity is limitless.

References

REF—1: AI was supposed to save coders time. It may be doing the opposite:

<https://www.scientificamerican.com/article/why-developers-using-ai-are-working-longer-hours/>

REF—2: Nvidia CEO delivers a blunt 7-word rebuttal on software stocks:

<https://www.thestreet.com/investing/nvidia-ceo-delivers-blunt-7-word-on-software-stocks>

REF-3: Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026:

<https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025>

REF-4: McKinsey forecasts up to \$5 trillion in agentic commerce sales by 2030:

<https://www.digitalcommerce360.com/2025/10/20/mckinsey-forecast-5-trillion-agentic-commerce-sales-2030/#:~:text=Those%20that%20hesitate%20could%20find,re%20redesigning%20what%20shopping%20means.%E2%80%9D>

REF-5: FutureScape 2026: Moving into the agentic future:

<https://www.idc.com/resource-center/blog/futurescape-2026-moving-into-the-agentic-future/>

REF-6: Beyond FTEs: reimagining SI asset strategy in the age of AI:

<https://www.everestgrp.com/blogs/beyond-ftes-reimagining-si-asset-strategy-in-the-age-of-ai/>



REF-7: The new growth playbook: capturing non-linear revenue growth through value-linked operating models:

<https://www.everestgrp.com/blogs/the-new-growth-playbook-capturing-non-linear-revenue-growth-through-value-linked-operating-models/>

REF-8: The Coming Battle For Share in SDLC Services:

<https://www.forbes.com/sites/peterbendorsamuel/2026/04/14/the-coming-battle-for-share-in-sdlc-services/>

The Authors

You can contact Sash Purohit , Chief Architect & SVP Engineering at sash.p@7t.ai and Sreedhar Kajeepeta , Advisory CTO at sreedhar.p@7t.ai



www.7t.ai

**Houston
Regional Office**

1334 Brittmooore Road
Suite D Houston,
Texas 77043
+1 (832) 632-4869

**Dallas Office
Headquarters**

16803 Dallas Pkwy
Suite 300 Addison,
Texas 75001
+1 (214) 299-5100

**Charlotte
Regional Office**

2115 Rexford Road
Suite #570 Charlotte,
NC 28211
+1 (980) 350-5100

