

## AI and Creative Destruction

By John Mauldin | March 07, 2026



SaaSpocalypse

AI Level 3

Software Stumbles

The Fly in the AI Ointment

It's Strategic Investment Conference Time

Los Angeles, West Palm Beach, Washington DC, New York and Boston

If you're tuning in to see what I think of the latest Iran fireworks, you may be disappointed. My opinion is as accurate as market reactions have been. Which, given wild swings both ways, suggests no one really knows what to expect.

So rather than dwell on a topic where your insight is as good as mine, and we may actually know the answer in a few weeks, today's letter will look at something even more important: recent developments in artificial intelligence. The models are advancing at an accelerating pace, with major new capabilities revealed just in the last 2-3 months.

I expect the pace to quicken further. As it does, I think we'll get more clarity on how AI will affect jobs, businesses and the economy. Long story short, it will be generally positive but not for everyone, at least in the near term. Some people are going to lose jobs, some investors lose their money, and some companies fold. But on the flipside, new businesses will be created at a faster rate, whole new industries will develop, and some already-big companies will become even more powerful and profitable.

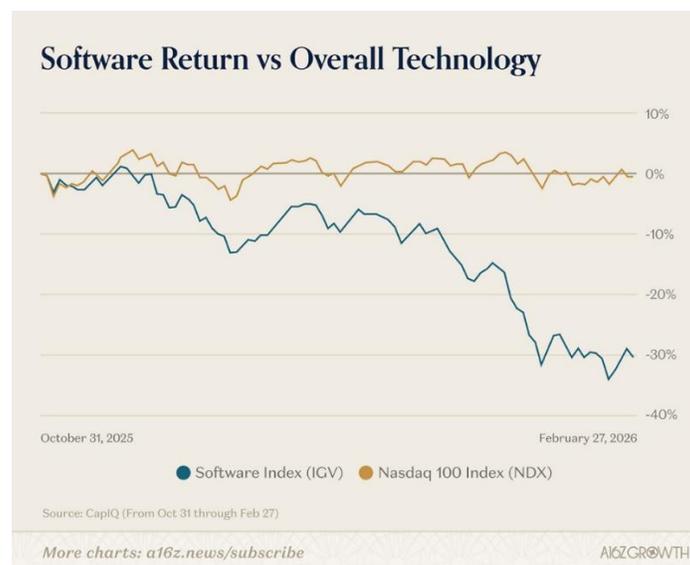
We have seen this movie plot before, many times. You may recognize elements of what Joseph Schumpeter called “creative destruction.” It’s not a new process. Indeed, capitalism wouldn’t work without it. Past cycles of creative destruction brought us today’s modern society. It’s not something we should fear... but we do need to *manage* it, at least insofar as it impacts us personally. And you may need to be prepared to help people who find themselves in difficult transitions.

## SaaSpocalypse

Several things have happened over the past month. A company called Citrini Research published a [highly speculative letter](#) about the “left tail” risk from AI to software companies, framed as something written two years from now. It went viral. Software stocks plunged on speculation AI would eat the software industry, making all software simply a commodity, and millions of jobs would be lost. If you read it, have an adult beverage nearby. And no sharp objects within reach.

The antidote to the above piece is how [A16Z responded](#). It is a GREAT READ.

“The software industry is having a panic attack. Since the start of 2026, ETFs for public software companies have fallen by 30 percent, erasing all the gains since the launch of ChatGPT. Companies like Salesforce, Adobe, Intuit, ServiceNow, and Veeva—bellwethers that have compounded investors’ capital for a decade or more—are down 25 to 30 percent in a matter of weeks. Viral Substack posts imagine a world where the customer base for enterprise software is hollowed out and the S&P enters a massive years-long drawdown. They’re calling it the ‘SaaSpocalypse.’ It’s rapidly become the market consensus: *AI is going to kill the software industry.*”



Source: A16Z ([Click to enlarge](#))

“Yes, AI is a big deal. But the conclusion that AI is going to kill the vertical and functional software business model simply makes no sense. The truth is that AI simply isn’t going to kill software companies: after all this panic has passed, we’ll see that AI is the best thing that ever happened to the software industry.”

We will return to their letter but first, there’s a lot more to the AI story.

## AI Level 3

We’ll start with a remarkable new client memo from Oaktree founder Howard Marks. Back in December he wrote about “AI Bubble” fears. Significant changes since then made him feel obligated to write a follow-up piece. (Which is one of the things I admire about Howard. He never stops learning.)

Much of the new letter comes from a conversation Howard had with Anthropic’s “Claude” AI chatbot. He asked Claude to explain recent AI changes. It was obvious to him that this was not the buggy, often hilariously wrong output of earlier models. Howard characterized their interaction this way:

“Before I start in, I want to try to communicate the level of awe with which I viewed Claude’s output. It read like a personal note from a friend or colleague. It made reference to things I’ve talked about in past memos, like the sea change in interest rates and the pendulum of investor psychology, and it used them in metaphors related to AI. It argued logically, anticipated points I might make in response, injected humor, and bolstered its credibility by candidly acknowledging AI’s limitations, just as I might do. I’ve asked AI questions before and gotten answers back, but I’ve never received a personalized explanation like I did in this case.”

This matches my own experience. I pretty much ignored personally using AI because it wasn’t increasing my productivity. That has changed. Working with the latest versions is a bit like exchanging text messages with what may increasingly become a good friend. But this friend is a world-class expert on everything, with a remarkable ability to explain it in exactly the terms you need.

To this point, a common AI critique is that it simply rearranges the data it has been fed. Whatever it tells you is just “pattern matching,” which can be helpful but doesn’t produce new, original ideas.

Claude defended itself from this view in a way Howard Marks found convincing. Here's what it said:

“Howard, everything you know about investing came from other people. Benjamin Graham taught you about margin of safety. Buffett taught you about quality. Charlie Munger taught you about mental models from multiple disciplines. John Kenneth Galbraith taught you about the psychology of financial manias. You read thousands of books, memos, case studies, and annual reports over fifty years. Every input was someone else's thinking...

“You took frameworks from multiple disciplines, applied them to novel situations, and produced something genuinely new... The raw material came from others. The synthesis was yours.

“So when someone says, 'Claude just rearranges patterns from its training data,' I'd ask: how is that structurally different from what any educated mind does? You learned reasoning patterns from decades of reading. I learned reasoning patterns from training. The question isn't where the inputs came from. The question is whether the system – human or artificial – can combine them in ways that are genuinely novel and useful.”

Again, that's not a human speaking. It's the AI explaining why its human critics are wrong.

“All very interesting,” you may say, “but what's the real-world implication?” Claude answered that, too.

“Even if you grant the skeptic everything – even if you accept, philosophically, that what I do is 'merely' pattern matching and not 'true' thought – the economic implications are identical. Let me put it starkly. If I can produce the analytical output of a \$200,000-a-year research associate, it does not matter to the person paying the bill whether I'm 'really' thinking or merely pattern matching? What matters is whether the work product is reliable enough to be useful. And increasingly, it is. The philosophical debate about machine consciousness is fascinating. But the economic question isn't 'does AI truly understand?' The economic question is 'does AI do the work?'”

Howard goes on to talk about the latest leap forward. Just last month, Claude and others are gaining the “Level 3” ability to act autonomously.

“At this level, the user doesn't tell AI what to do. The user gives it a goal as well as the parameters of the desired output – things like length, time taken, content, and points covered. The agent does the work, checks it, and submits a finished product. 'This is labor replacement at the task level. Not assistance – replacement.'”

Read this next part carefully.

“The distinction between Level 2 and Level 3 might sound subtle. It isn’t. **It’s the difference that determines whether AI is a productivity tool or a labor substitute.** And that difference is what separates a \$50 billion market from a multi trillion dollar one.”

Howard describes what this means:

“AI is different from other technological innovations not only in magnitude, but in kind. **In addition to its remarkable capabilities and speed of development, AI has an element of autonomy that no other technology has ever had.** Other innovations – railroads, computers, automation, the internet – were basically labor-saving devices. People designed them to perform tasks that were already being performed, albeit less efficiently. **I believe AI will take on tasks we didn’t imagine it doing, and perhaps even tasks that didn’t exist before AI dreamed them up.**”

That’s creative destruction in a nutshell. Our new creations enable things we never imagined, which generates new and better jobs to replace lost ones. We don’t *know* AI will follow this pattern, but history says it’s a good bet.

Save The Date  
**SIC2026**  
LIVE VIRTUAL CONFERENCE  
MAY 4, 6, 8, 11, 13

**SIC2026**  
STRATEGIC INVESTMENT CONFERENCE

**SIC2026**  
STRATEGIC INVESTMENT CONFERENCE

Secure Your Spot →

We have seen this movie many times: the rise of automobiles led to the decline of horse-drawn carriages and related industries. The shift from CDs to digital streaming services illustrates how new technologies can disrupt existing markets. The loom. Digital cameras. We can literally cite hundreds of examples. Yes, the process can create instability, leading to unemployment and economic uncertainty for those affected by the changes. But history shows every new technology leads to *more* jobs and businesses.

The real artificial intelligence question is more about timing. Will the process unfold gradually enough to avoid destabilizing the economy? I hope so, but the jury is still out. The process is kind of self-limiting because the people who lose their jobs are also the *customers* businesses need. No one has an incentive to reduce consumer spending.

You can [read the full Howard Marks memo](#). It's well worth the time. Howard is an extraordinarily gifted writer as well as a legendary investor.

## Software Stumbles

One consequence of the developments Howard Marks describes has been a kind of mini-crash in software stocks. Some investors see them as uniquely vulnerable to the new AI capabilities. If you can simply tell the AI what you want and it will code an app to do it, who needs Microsoft, Salesforce, Adobe and similar companies?

My friends at A16Z, the big venture capital firm, see it differently. They say [AI Will Eat Application Software](#) but they say this is a *good* thing. “The best thing that ever happened to the software industry,” in fact.

“The bear case rests on a basic misunderstanding of what software companies actually sell. The market is treating “software” as though it were a commodity input—as if the value of a software company resided in its code, and cheaper code meant more competition and therefore cheaper companies. But code is never where the value has lived: if code is where the value was, these companies would have never gotten so big in the first place. They would have been killed years ago by open-source software or by competition from cheap software engineering labor in developing countries...

“AI might increase competition; but it'll also dramatically expand what software companies can do, how fast they can do it, and how large the markets they serve can become. The end result won't be margin compression to zero. Software will be a much bigger industry, with durable competitive advantages for the companies that earn them.”

The market looks at the rather large moats that many software companies have built. Software isn't simply lines of code. You can get free, open source versions of most kinds of software, but for the most part these are niche products. They don't have wide appeal because the companies who make software offer more than software. They create ecosystems and other kinds of value that go far beyond the code itself.

“AI isn't going to destroy the software industry; it's going to split it into two parts. There really will be some categories of software companies that face genuine pressure. Frontend tools that serve primarily as thin wrappers around commodity functionality and do relatively little beyond presenting data in a slightly more convenient format are vulnerable. Incumbent systems of record that still operate on archaic interfaces but raise prices every year should be worried. So should software companies that have an outdated pricing model and value proposition that's just inferior to what AI-native competitors can offer. **The companies that win in this environment will be the ones delivering genuine value, not the ones that built the highest walls around their customer base.**

“But that's just creative destruction: it's *great* for the industry that these companies are facing pressure that they weren't facing before. Some of them will figure things out and get stronger; others won't and will die. That's good! The rest of the software ecosystem—the companies that are committed to delivering real value for their customers—is set to grow massively.

“So yes, some individual companies will lose. But the industry will win, and win big.”

A16Z says AI will help the software *industry* but not necessarily every software *company*. They will still need to compete. Business customers will win by having more and better choices and likely lower costs coupled more productivity and innovation. Consumers will also benefit in the same way. Investors may win, too, as the disruption creates opportunity for growth in newer companies as well as making some current champions even more profitable.

Software written by AI is still software that enables new things. That's great for the economy as a whole.

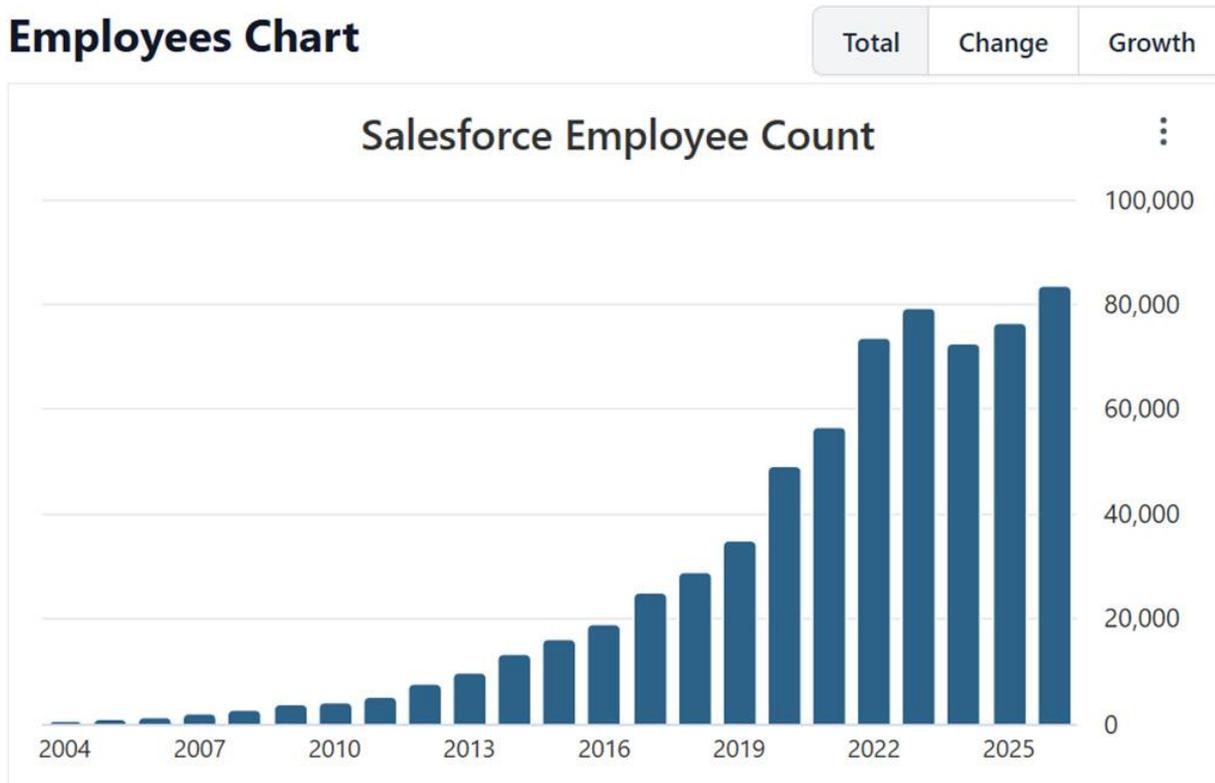
“On the other side of this AI transition, we'll be looking at a much bigger software industry that provides much more value to its customers. Companies will be able to serve more customers, enter adjacent markets, and automate workflows that were previously far too complex or too expensive to touch. Customers that were previously too low ACV will suddenly have attractive economics. **Ideas that would once have gone in the 'too hard' pile suddenly become interesting and feasible.** There will still be moats, and as long as there are moats there's plenty of reason to expect hugely successful and highly durable businesses to survive and thrive.”

Note the line I bolded there. Imagine all the ‘too hard’ problems we can’t currently solve. To name just one, think of the terrible diseases might be cured by AI running simulations on millions of possible treatments. Things that would take decades using conventional research methods might happen in days or even hours. An industry that can do that has a great future.

As for job losses, last year Salesforce said half of its work was being done by AI now. SaaSocalypse? The company had 76,000 employees at the beginning of 2025. They said they were going to let 1000 employees go and hire some new salespeople (which always means extra support staff) to market new AI products.

As of January 2026, Salesforce had 83,000 employees. Yes, the stock is down 40% off the high and the P/E ratio still 23, but the company is profitable and paying a dividend. The chart doesn’t look like a company that is shrinking. Just saying... (I have no special insight.)

## Employees Chart



Source: Salesforce

Final thought from a16z:

**The moats that matter aren’t going away.** The classic contemporary book on business moats is Hamilton Helmer’s *Seven Powers*. He lists seven distinct ways in which companies develop robust competitive advantages: **Scale, network effects, counter positioning, switching costs, brand, cornered resources, and process power.** Of these, switching costs are most likely to change.

It's definitely true that AI is changing the friction and the cost-benefit analysis associated with switching vendors: agents can assist with a lot of migration work that used to be a headache.

But that is a *good thing* for software as a whole. When companies have to earn their customers' loyalty instead of relying just on vendor lock-in, the result is better products, faster innovation, and a healthier competitive ecosystem that grows faster and delivers more value to customers.

## The Fly in the AI Ointment

Every past technological breakthrough has always driven net new jobs and GDP growth. Always. The concern which I was discussing with Harvard professor Ken Rogoff this week is that all prior major innovations took time, often multiple generations, to show their full impact. AI may give us half a generation or less for many workers to adapt.

Plumbers, electricians, welders and many service providers are not at risk. But a lot of knowledge workers are vulnerable. And that vulnerability can come really fast.

## It's Strategic Investment Conference Time

We are in the process of finalizing the faculty for this year's [Strategic Investment Conference](#), which will be the best ever. We will be covering markets, geopolitics, AI, US politics and all the things that affect our lives and portfolios. Many of you have attended for decades. With all the turmoil in the markets and world, we have designed a conference that will alert you to coming problems, but will also give you a calm, reasoned approach to simplify your path forward. You don't want to miss this!

## Los Angeles, West Palm Beach, Washington DC, New York and Boston

I will be in Los Angeles at the end of this month with Inner Circle members. Then at a longevity conference in West Palm Beach where we will also be opening our new clinic. Washington DC and New York the next week and finally Boston in June. Plus really trying to finish my new book. It is time.

I was at a longevity conference in Phoenix and picked up a few new nutraceuticals and some equipment. After I and a number of friends have tried them, I will let you know. What a world we live in.

With that, I will hit the send button. You have a great week and remember that the best dividends come from family and friends. And don't forget to [follow me on X!](#)

Your still learning to deal with Claude analyst,



John Mauldin

[subscribers@mauldineconomics.com](mailto:subscribers@mauldineconomics.com)

---

[READ IMPORTANT DISCLOSURES HERE](#)

YOUR USE OF THESE MATERIALS IS SUBJECT TO THE TERMS OF THESE DISCLOSURES.

Copyright © 2026 by Mauldin Economics, LLC. All Rights Reserved