


THE STATE OF DEVELOPMENT 2026

How engineering teams are using AI agents

We surveyed **554 engineers and engineering leaders in the US and UK**



[Read this online →](#)



CAPTAIN'S LOG

The year is 02026.
Where are all the
AI-using engineers?

CAPTAIN'S LOG

They're at the frontier. And they left messages behind. But what of it is true? And what's not?

CAPTAIN'S LOG

Command needs
us to find out.

Initiate survey n=554



What's inside?

- 01 What are engineering teams using AI for?
- 02 What effect have AI agents had on their organization?
- 03 What's holding teams back from using AI agents more?
- 04 How are they feeling about it?
- 05 What does the agentic near future hold?

Survey methodology

We're on the trail of ...

Mostly engineers with 6-15 years of coding experience (58.9%)

- > Engineer / AI engineer (25.6%), VP / director of IT (13.9%), data engineer (11.9%), devops/platform engineer, architect, data scientist, CEO / management

Working at mid-sized companies, over one-third of which are in software

- > 251-1000 employees (29.2%), 51-250 employees (22.7%), 1,001-5000 employees (18.8%)
- > Software (37.5%), manufacturing (9.7%), professional services (9.7%)

Located in the US (2/3) and UK/EMEA (1/3)

Definitions and caveats

- + AI agent = An LLM instance capable of taking multi-step actions on its own.
- + When we discuss those who are “successful at using AI agents,” that data is self-reported.
- + Between April 29 to May 25, 2026, we commissioned the survey provider Qualtrics to solicit responses from 650 respondents currently using AI agents following strict criteria around role and location. We removed all responses we considered low quality and were left with 554 respondents.

Key findings

One group of engineering teams is using AI much more effectively. Throughout this log, we show how they fare by comparison. But generally, we found:

A 70.8% leap in AI agent use: 80.8% use agents daily, up from 47.3% a year ago (1.2)

- > 21.8% use AI agents continuously (1.2)
- > 21.8% say agents are core to how they ship (1.1)
- > 25.5% still just using them as assistants (1.1)
- > Actual use and desired use don't diverge much (1.8)

Respondents run a median of 5.0 agents (1.5)

- > The mean is much higher: 10.7 agents (1.5)
- > Just 2.2% are running 51+ agents (1.5)

Top AI agent uses: #1 writing code, #2 testing code, #3 analyzing (1.6)

91.1% say agents have 'improved' or 'revolutionized' their productivity (2.1)

- > 51.3% go from AI prototype to production-ready code in hours or faster (2.2)
- > The successful aren't substantially faster, just 1.2x (2.2)
- > 85.5% trust or somewhat trust the outputs (2.4)
- > But also: 41.1% encounter issues daily or more and 9.0% say "continuously" (2.6)

Top blockers to using AI agents more: #1 tracking state, #2 debugging, #3 managing costs (3.2)

- > 79.8% say token and compute costs are a limiting factor (3.4)
- > Top places they troubleshoot: #1 YouTube, #2 AI tools, #3 private Discord or Slack group (2.7)

The SaaSpocalypse is real: 92.3% have tried to rebuild a software they used to buy (2.9)

Key findings

84.5% believe they are better than their competitors at using AI agents

- > This is somewhat driven by executives, who are 35.1% more likely than their teams to say this (2.5)

They feel relaxed: 77.5% are more optimistic about the future of their job than a year ago (4.5)

- > 44.6% are less stressed than one year ago (the successful are much less stressed) (4.1)
- > Many say that actually, agents aren't the primary source of stress (4.3)
- > People are optimistic about a future with more agents, but also worried (4.4)
- > 56.7% think it'll be harder for junior people to find jobs (4.6)
- > 45.5% think it'll be harder for senior people to find jobs (4.7)
- > Yet only 26.4% of companies say they are slowing or stopping their hiring (4.8)
- > A large number think AI will be 'positive' for the climate (4.11)
- > 8.8% avoid using agents more because of societal or climate concerns (3.2)

Key themes

Engineering teams that use AI agents are a lot less worried than we expected

AI agents are more useful than not. Many are excited by them. These respondents sense that agentification is inevitable. Some worry how agents might erode their skills, freedom, and the environment. But overall, they'd caution the public to not worry as much. Engineers have many reasons to feel stressed. Agents are but one.

The code-gen step has collapsed

When 51.3% say they now go from prototype to production-ready code in hours or faster, and 26.9% say minutes or faster, ideas are quick to execute. But this means it's more important to make the right decisions. Which means engineering fundamentals matter more.

They are poised

Engineers think agents are getting plenty of societal attention. They aren't stressed and don't think you should be worried about them because AI agents aren't the primary cause of stress — the job creates plenty of that on its own.

They are more likely to be less stressed than one year ago

- + 32.5% are more stressed
- + 22.9% are the same or "other"
- + 44.6% are less stressed

Key themes

They are excited

77.5% say they are “very optimistic” about their role. AI agents have made the majority more efficient.

Though when asked to reflect, they worry

The free text responses reveal a good deal more anxiety than comes across in the multiple-choice questions. The more we ask them to reflect, the more the worry unfurls.

- + “Cautiously optimistic”
- + “Nice for me”
- + “High impact to human evolution”
- + “As long as humans remain responsible for validation”
- + “I see their potential but worry about over-reliance and skill erosion”
- + “Concerned about maintaining critical thinking skills”
- + “Overrun with AI”
- + “I hope it’s not a riot fire by then seriously”
- + “More hostile to actual people”

Some fear ever-moving goals

Respondents worry that no matter how much more efficient they grow, leadership will just keep raising expectations. People at successful companies are less likely to say this.

They risk slipping into greater isolation

Before asking a teammate for help with agents, people will look to YouTube, LLMs, Discord, chat groups, academic papers, GitHub, and X (Twitter).

They worry about the loss of skills and freedom

In free text responses, they worry about what’s lost when they collectively rely on automated tooling. Many anticipate a near future with fewer jobs overall, including for developers.

- + “I have mixed feelings — they promise efficiency gains, but also bring challenges like skill adaptation and the risk of over-reliance”
- + “It’s going to be a lot harder to find work in this sector”

Key themes

The successful are far more confident


While everyone agrees agents have helped, the successful are more certain. They use more tools, have less interest in changing course, know to search for help in more places, and are 2.2x as likely to say they are less stressed than one year ago.

They predict

- + “Programming will be more abstract”
- + “The software engineer role will become an agent fleet orchestrator”
- + “It will require more and more creative minds”
- + “It will be tremendously helpful in many aspects of life”
- + “It will be furious — They will improve everything”

”

Let's review the full results




01 What are they using AI for?

49.1% say agents are “in production” or “core to how they ship”

The AI-using cohort has traveled deeper into agent territory than we thought.

Chart 1.1

How would you characterize your team's use of AI agents?



Frequent AI agent use jumped 70.8% in the last year (daily or more)

For comparison, it took Agile 11-15 years to reach mainstream. Here, we've seen AI agent usage increase from 47.3% saying “daily or more” to 80.8%.

Chart 1.2

How frequently do you use AI agents? (one year ago vs. now)






Chart 1.3
How frequently do you use AI agents? (one year ago vs. now); visualized differently



The successful were 1.5x more likely to use agents daily or more (85.5% vs 55.8%) and 2.3x more likely to use agents continuously (23.9% vs 10.4%)

Chart 1.4
How often did you use AI agents one year ago vs. today?




The median respondent is using 5.0 agents

Though on average, they use 10.7. That 2.1x difference between the mean and median indicates that there are a few very large outliers — individuals or teams running legions of agents, such as those who replied “256.”

The successful run 1.3x more agents on average (11.1 vs 8.8).

Chart 1.5
How many agents are your team running, approximately?



	Mean	Median	Min	Max
The successful	11	5	1	256
Everyone else	9	4	1	123

The primary use cases are writing code, testing code, analyzing, generating ideas, and debugging


After that, there’s a modest dropoff to the bulk of lesser use cases — though none was used by less than 14.8% of all respondents.

The successful are much more likely to use AI agents for technical design, security, and support. They also use AI agents for more use cases.

The successful are more likely to use AI for:

- + 2.7x for technical design (28.0% vs 10.5%)
- + 2.6x for security (33.5% vs 12.8%)
- + 2.2x for customer support (28.0% vs 12.8%)
- + 1.6x for writing code (66.7% vs 40.7%)
- + 1.5x for researching knowledge bases (45.9% vs 31.4%)

Chart 1.6
What are you and/or your team using AI agents for?




Note: Respondents could select multiple options

Do teams want to use AI differently? Not really

We asked engineering teams what they want to be using AI agents for, to compare, and were surprised to find little difference. Is this because the future is unclear? Or that roles and uses are changing too quickly to see ahead? Or that they simply already know what agents are good for? (Write-in responses suggest the final option is likely.)

Chart 1.7

What are you and/or your team using AI agents for? (Current use vs desired use) Note: Respondents could select multiple options



If we zoom in, the successful are somewhat more certain that they’re using AI agents the way they want — there’s less of a gap between their present usage and desired usage.

Chart 1.8

What are you and/or your team using AI agents for vs how do you want to? (Successful vs others)

Note: Respondents could select multiple options




The top-used AI tools: #1 ChatGPT or API, #2 Microsoft Copilot, #3 Gemini, and #4 Claude

OpenAI has maintained its first-mover advantage and first place. It also has more offerings. If we combine the model offerings from OpenAI and include services by others like Microsoft’s Copilot that run on OpenAI’s tools, it accounts for most of the list.

Copilot and Gemini’s success are likely due to distribution — those are widely embedded in work tools, and defaults can be powerful. Claude is a close fourth, just above ChatGPT Codex.

The average team is using 3.5 of these tools.

Chart 1.9 Which of the following AI tools are you and/or your team using? Note: Respondents could select multiple options



The successful are using 1.3x more tools than everyone else (3.6 tools vs 2.7 tools), suggesting that when you get good at orchestration, you're either able to use more tools — or have to.


The successful are more likely to use:

- + AWS Bedrock (23.1% vs 8.1%)
- + Azure OpenAI Service (30.3% vs 18.6%)
- + ChatGPT Codex (34.4% vs 24.4%)
- + Claude (41.2% vs 31.4%)
- + GitHub Copilot (32.7% vs 15.1%)
- + Google Gemini or Vertex AI (44.2% vs 29.1%)
- + LangChain or LangGraph (7.3% vs 2.3%)

Temporal users are 3.1x more likely to use open source models (37.0% vs 12.4%). They also use more tools on average (5.0 versus 3.5).

Chart 1.10

Which of the following AI tools are you and/or your team using? (Successful vs others)




Note: Respondents could select multiple options

Agent frameworks: Again, OpenAI products were the most commonly selected

Compared to last year, Microsoft’s agent framework made gains, as did AWS Strands and CrewAI.

The successful are 1.7x more likely to be building agents on OpenAI Agents SDK (51.1% vs 30.2%) and are 1.9x more likely to use Temporal (9.0% vs 4.7%).

Chart 1.11
What frameworks are you or your team using to build agents?



Note: Respondents could select multiple options

Orchestration: We’re seeing falling interest in AWS Step Functions, LangGraph, and Airflow

Companies are only modestly looking to change how they orchestrate their agents, with declining interest in AWS Step Functions and LangGraph, and rising interest in Camunda, Inngest, and Prefect. Temporal, Dagster, and Orkes held roughly constant.

The successful show less of a desire to change how they orchestrate than everyone else. If it’s not broken, don’t fix it.

Chart 1.12


How do you currently orchestrate your workflows vs how you plan to?



Note: Respondents could select multiple options

Chart 1.13

How do you currently orchestrate your workflows vs how you plan to? (Successful vs others)



Note: Respondents could select multiple options

■ AWS Step Functions ■ Apache Airflow ■ Camunda ■ Dagster ■ Inngest ■ LangGraph ■ Orkes Conductor ■ Prefect ■ Restate ■ Temporal


02 What effect have AI agents had on you and your organization?

91.1% say AI agents have 'improved' or 'revolutionized' their productivity

Nearly everyone agrees that agents make developers more efficient. Only 4.0% said they had no impact and just 1.8% said agents had worsened their productivity.

Chart 2.1

How have AI agents altered your own productivity?




51.3% go from AI prototype to production-ready code in hours or faster

77.1% say they get there in days or faster. This is a striking finding given that 33.4% of our dataset were respondents from companies with 1,001+ employees, and 14.6% of companies had 5,001+ employees. So this isn't just a function of startups.

If we look at just that cohort of respondents from companies with 5,001+ employees, they fall into two categories: The very fast and the very slow. Most use agents to deploy at startup speeds, while one-fifth take "months" (18.5% replied that way, versus 5.5% of all others).

Chart 2.2

How long does it typically take to get from a working AI-generated prototype to production-ready code?



The successful are only 1.2x faster

A number of engineers told us that their biggest bottleneck is just plain old time to think. That shows up in several of the results where the successful aren’t substantially faster (78.9% daily or faster vs 67.4%). Yet their outcomes are better. AI speeds everyone up equally. Some make better use of it.

Positive

- “It helps greatly with complex debugging issues”
- “It has freed time to concentrate on important issues”
- “AI has made it easier to get larger loads of work completed in a shorter period of time”
- “It is very productive to use AI”
- “It increases my organization productivity by 25%”
- “Saves me 2-3 hours per day”


Neutral

- “I think improved in some ways but worsened in others as I have to check things regardless so it evens out”
- “I wish there was just one centralized AI agent. We really don't need more; companies aren't drastically improving”
- “I just don’t think it has any impact”

Negative

- “I think it has made me feel lazier”
- “They’re rooted in the efficiency drains I encounter most frequently”
- “When we lose morals our productivity goes down too”
- “I don’t like it”
- “Terrified”
- “It is scary”
- “Not looking forward to them”
- “I believe it will ultimately fail”
- “A bit concerned”
- “Very unsure, it could be great but also could cause an economic collapse”

Chart 2.3
Why did you respond that way about how agents impact your speed?




85.5% trust the outputs at least somewhat

Trust is high. One quarter (24.7%) completely trust the outputs. A majority (60.8%) at least somewhat trust. Only 4.0% distrust. These levels of confidence are surprising given all the public discourse around code slop and hallucinations.

The successful are far more trusting: They are 6.1x as likely to ‘completely trust’ (28.4% vs 4.7%) and 1.6x as likely to ‘somewhat trust’ (90.8% vs 57.0%).

Temporal users are 1.1x as likely to say they somewhat or completely trust their agents (95.7% vs 84.7%).

Chart 2.4
How much do you trust the accuracy of AI agent outputs today?




84.5% believe they are better at using agents than their competitors

A great majority of teams we polled think they’re in the 85th percentile of agent excellence which is, statistically, unlikely. (Possible, but unlikely.) We choose to interpret this favorably: There is a lot of misinformation out there about agent use. Most teams aren’t informed about what others are actually up to, hence the need for this report.

No surprise, though, some of this exuberance is driven by leadership: CEOs, management, VPs, and directors were 1.4x more likely to say their organization is ‘very successful’ (43.7% vs 32.4%).

Temporal users say they are 1.1x as likely to say they are successful at using agents than peers (91.3% vs 83.8%).

Chart 2.5
In your opinion, how successful is your company at using AI agents relative to peers?




Leadership is 35.1% more likely to say their organization is ‘very successful’ compared to the engineers.

41.1% encounter issues with AI agents daily or more

16.4% say they encounter issues hourly or more, which sounds punishing. Surprisingly, the number of agents teams deploy and the frequency of issues are not correlated. Teams running dozens of agents do not necessarily encounter more issues than those running five. Though, a caveat: Each team may define issues differently, issues can range in severity, and those who pay closer attention will notice more of the granular issues.

Successful teams are just as likely as others to experience issues “daily or more.” But expand that and you see that the successful are 1.7x more likely to experience issues “continuously” (more frequently) and everyone else is 1.3x more likely to experience issues daily (less frequently).

Chart 2.6
How often do you encounter issues with AI agents?




YouTube is the #1 place teams go for help troubleshooting agents

It is interesting that YouTube, a video site where viewers cannot as easily copy the code they see onscreen, beats GitHub and Stack Overflow, where they can. (It is possible to copy code from YouTube with Gemini, but less convenient. It sometimes pulls only fragments.) This suggests that AI agent developers would rather watch someone walk through the results onscreen than copy a pattern. Maybe this means the temporal aspect of running many agents makes it more important to understand the workflow than the exact code, which they can easily generate.

In future research, we are curious to know whether this suggests a trend toward isolation: Engineers will look at videos, Discords, chat groups, academic sources, GitHub, and X (Twitter) all before they talk to someone at their own organization. This sense of collaborative distance shows up in the written responses.

Also: 1.8% of people think they're in the smartest 1%. Talk about *accuracy*.

Chart 2.7
When you encounter an issue with an agent, what are all the places you'd consider turning to for guidance?



Note: Respondents could select multiple options

Engineers will look to videos, LLMs, Discords, chat groups, academic sources, GitHub, and X (Twitter) before they talk to someone at their own organization.


The successful are more resourceful troubleshooters: They search 4.1 places, where the others only search 3.5 places.

The successful are also:

- + 2x as likely to turn to an AI tool for help (40.4% vs 19.8%)
- + 1.9x as likely to turn to X (Twitter) (31.6% vs 16.3%)
- + More likely to turn to peers, LinkedIn, Stack Overflow, Reddit

Chart 2.8

When you encounter an issue with an agent, what are all the places you'd consider turning to for guidance?
(Successful vs others)



Note: Respondents could select multiple options

92.3% have tried building an app they'd previously have bought

25.6% said they succeeded in building an app they'd otherwise have paid for, and that it had a big impact. This suggests the so-called “SaaSocalypse” is real — teams are finding the tools to build new tools and aren’t coming back as frequently for resupply.


The successful are 8.5x more likely to have rebuilt something and said it had a big impact (29.7% vs 3.5%). Everyone else is 2.9x as likely to say they didn’t try or it didn’t work (17.0% vs 6.0%).

Temporal users are 1.6x more likely to say that they’re rebuilding rather than buying and it had a big impact (39.1% vs 24.4%).

For future research: How will devops teams adapt when there are far more custom/vibe-coded tools that don’t necessarily interoperate? Does all this self-generated tooling make it harder for developers to switch jobs? Does it mean they’ll require longer to onboard at the new one?

Chart 2.9

Has your team built (or are you planning to build) something internally that you would have previously bought a software for?




CAPTAIN'S LOG

We've almost caught up.

And...

CAPTAIN'S LOG

There they are. The
forward-most deployed
AI agent engineers.



CAPTAIN'S LOG

They're all converging.

03 What's holding teams back from using AI agents more?

Top blockers to using AI agents more: #1 tracking state, #2 debugging, #3 managing costs


In our survey, 35.7% said the most common blocker to using more AI agents is tracking state, followed by debugging (agents or otherwise) and managing costs (tokens or compute).

Successful teams had more responses per respondent, meaning they face a wider variety of issues — or at least are better aware of them. They were somewhat more likely to say, “tracking,” “retries,” “unrelated business logic,” and “infrastructure,” and somewhat less likely to say, “orchestrating,” “isolating,” or “managing costs.”

Chart 3.1

In general, what limits you or your team's productivity?

Note: Respondents could select multiple options



Where are the truly autonomous agents? 39.5% cite security concerns

The greatest blocker to launching self-managing agents is the same thing that has always blocked teams from doing more outsourcing: Engineers need certainty that their delegates have done their job correctly. The only way to know — beyond time-consuming testing and mathematical proofs — is to manually check.


Thus, even while 85.5% of respondents trust the agents' outputs at least somewhat, there is a difference between trusting enough to proceed and trusting enough to put your name on the commit.

Successful teams worry somewhat less about the unknowns, such as lack of reliability or invented results, unintended consequences, and unpredictable outputs.

Temporal users are 2.2x as worried about AGI (21.7% vs 10.4%). They are 1.4x less likely to be blocked by poor integrations with existing systems (15.2% vs 22.6%) and 2.5x more worried about the climate impacts of AI agents (19.6% vs 7.9%).


There is a difference between trusting enough to proceed and trusting enough to put your name on the commit.

Chart 3.2
What's blocking you from truly autonomous agents?



Asked to elaborate, respondents volunteer more skepticism and uncertainty. They acknowledge more risks and drawbacks.

Chart 3.3
Why did you select those blockers?




Write-in responses

- “Our network is quite locked down and not all apps integrate with AI, so we can't use it for everything”
- “The technology is still new, and we must use it with caution”
- “They aren't very reliable as they can be very varying and have stochastic reasoning, intolerable risks of making unapproved mistakes in production”
- “I don't want to make people redundant yet”
- “Isn't reliable and safe for us to use on main production”
- “The learning curve is too high”
- “I don’t want to sign off on hallucinated content from the AI”
- “We worry this may make consumers feel different about our brand”
- “We have a lot of security rules at our workplace”
- “These blockers align with my manufacturing engineering work, where security, compliance, and the overhead of non-core tasks like debugging and boilerplate code slow down AI adoption and progress”
- “Because sometimes AI can be more hassle than is needed”
- “The ‘heads of’ are scared of potential data leaks which can stunt growth”
- “AI is also difficult to monitor”
- “Personal values and climate awareness drive my choices”

79.8% say token and compute costs limit their progress

Companies are realizing with chagrin that in some cases, tokens can be quite expensive. The successful are 1.3x more likely to say cost is at least somewhat of a factor (83.1% vs 65.7%). Everyone else is more likely to say they’re unsure (22.1% vs 7.3%).

Chart 3.4
Is the cost of running AI agents a meaningful factor in your decision to use AI agents?




04 How are they feeling about it?

44.6% of respondents were less stressed than one year ago

We expected panic. Instead, we found a spacey calm. Even at the frontier of progress, life goes on. There's gravity. People need to eat. Respondents are more likely to say they are less stressed compared to a year ago.

Chart 4.1


What are your current stress levels compared to one year ago?



Teams that are successful are 2.2x more likely to say they are feeling less stressed than one year ago (48.7% vs 22.1%), and are less likely to say they are more stressed (30.4% vs 44.2%). This is because everyone else is so likely to say that nothing has changed for them year-over-year. Whereas for the successful, things are changing one way or the other.

The more stressed the team, the more issues they encounter: 52.2% of those who are more stressed encounter agent issues daily, versus 32.4% of those who are less stressed.


Chart 4.2
Stress levels of the successful vs others



Many say agents do not factor into their stress

At the frontier, software engineering is difficult and agents aren’t the biggest factor. Those in technology can sometimes forget that the new thing isn’t everything, and that managing their time, bureaucracy, organizational disorder, tooling, meetings, and testing are still the bigger issues. And perhaps will always be.

Chart 4.3
How do AI agents factor into your answer about stress, if at all?



Write-in responses

- “I’m much less stressed because the services we use are very reliable”
- “Not really at all, I’m stressed because of the workload and uncertainty”
- “Automation eases, but overnight it intensifies”
- “Because the AI standards are not yet fully developed, they don’t touch upon my core work; they only help me with repetitive tasks”
- “Job in general is stressful”
- “They play a part but not a majority”
- “I don’t think they impact it much, outside of maybe people thinking AI content is real”
- “I just get more work no matter what”
- “We are judged based on how much AI we use”
- “It is always stressful trying to produce items while also learning how to do it at the same time”
- “Sometimes, I feel hopeless that I will lose my job to AI”

People are optimistic about a future with more agents, but also worried

Chart 4.4
How do you feel about a future with more AI agents?



Write-in responses

- “Cautiously optimistic”
- “Nice for me”
- “High impact to human evolution”
- “I’m optimistic about a future with more AI agents as long as humans remain responsible for validation, decision-making, and risk management.”
- “I LIKE IT”
- “I see them as a tool for progress, as long as we implement them responsibly and inclusively across industries”
- “I see their potential but worry about over-reliance and skill erosion”
- “Excited about the efficiency gains, but concerned about maintaining critical thinking skills”
- “I have mixed feelings — they promise efficiency gains, but also bring challenges like skill adaptation and the risk of over-reliance”

77.5% are more optimistic about the future of their role than a year ago

The successful are 8.4x more likely to say they are “much more optimistic” (39.1% vs 4.7%). Though both groups are equally likely to say “slightly optimistic.”

The successful are 6.4x as likely to say they are much more optimistic about the future for their peers, too (37.2% vs 5.8%).

Temporal users are 1.2x more likely to say they are optimistic about the future of their role (89.1% vs 76.4%).

Chart 4.5
How optimistic do you feel about the future of your role compared to one year ago?




Chart 4.6
How do you think AI agents will affect the job prospects for junior people?






Chart 4.7
How do you think AI agents will affect the job prospects for senior people?



Despite the fear for others' jobs, only 26.4% said their companies were 'stopping' or 'slowing' their hiring

Everyone thinks the prospects for hiring are worsening. But their own companies tell a somewhat less pessimistic story: Most are still hiring, and many are prioritizing AI agent experience. The spaceship, it seems, is well-insulated.

Chart 4.8
How optimistic do you feel about the future of your role compared to one year ago?




The successful are more likely to be hiring, and hiring evenly: They are equally likely to be hiring for senior as for junior roles. They are also 1.8x more likely to be hiring for “AI agent” experience (47.9% vs 26.7%).

Everyone else, on the other hand, is 1.7x as likely to be hiring for junior roles as they are senior roles, or are waiting to see how things play out: They are 2.7x more likely to say AI agents have had “no impact” on their hiring plans (20.9% vs 7.7%).

Chart 4.9

How are agents influencing how your team is structured? (Successful vs others)




29.6% think agents are getting too much societal attention

Everybody — 98.5% of respondents — had an opinion on this. The successful are more likely to feel the societal attention is correctly calibrated, and less likely to have “no opinion.”

Chart 4.10

Do you think AI agents are getting the right amount of societal attention?



A large number of developers think AI agents will have a positive ecological impact

Will AI usher in an era of scientific discoveries that solve climate change? Or will it just make things worse in the fruitless pursuit? Developers lean toward the former.

Though, recall that 8.8% cited ‘societal or climate’ worries as a reason for not using AI agents more. Those who felt that agents would have a negative climate impact tended to fall into the everyone else cohort, which is not commentary on our part, just an output from the survey.

Chart 4.11
How do you feel about AI's ecological impact?



Looking 5-10 years out, teams think AI reliance is inevitable

In a world with more AI agents, they largely agree that engineers will act more and more like “natural language orchestrators” over fleets of agents that handle code generation, integration, and maintenance.

Themes among these responses:

- + “It will invite in human creativity”
- + “Code will automatically integrate”
- + “Programming will be instantaneous”
- + “The “Developer” and “AI trainer” distinction will be blurry”
- + “There will be many 1-person “teams”
- + “Code-gen will become far more inclusive to the non-technical”
- + “There will be fewer jobs, including for engineers”

Chart 4.12

What do you predict 5-10 years out?



Write-in responses

- “Overrun with AI”
- “Worse than now as there will be less access to AI”
- “It will require more and more creative minds as we reach higher levels of AI”
- “In the near future, everything might be computer or technology based thereby making everything easier for human adaptation”
- “It's going to be a lot harder to find work in this sector”
- “Programming will be more abstract”
- “Heavily agentic”
- “Fighting legal battles lol”
- “Flying cars”
- “Scary”
- “The software engineer role will fully transition into an agent fleet orchestrator”
- “All AI-driven system architecture and natural language orchestration”
- “IT WILL BE WIDELY ACCEPTABLE”
- “So far from what we have now that it won't be believable”
- “More hostile to actual people”
- “I think things will get more complicated because we have unethical people in society”
- “I hope it's not a riot fire by then seriously”
- “I'm a little nervous about all that”
- “It will have made its money and destroyed lives and land in the process”
- “It will be outstanding and tremendously helpful in many aspects of life”
- “More organized and clean”
- “It's not a big deal”
- “It will be more inclusive of people who may not be passionate about coding or computer science as a whole”
- “We're gonna have a lot of overweight people from laziness”
- “It will be furious; they will improve everything”

Their advice to us in the present? Many say “chill out”

Answers of course range, and include responses like “panic,” “boycott,” “fearmongering,” and “worry,” but the most common responses were (semantically): relax, adapt, learn, chill, calm, focus, use, embrace, learning, and prioritize.

Chart 4.13
From what I'm seeing, everyone just needs to -----.



Write-in responses

- “Chill”
- “Invest in learning”
- “Take a step back and look at what is happening”
- “Have fun and relax”
- “Wise up”
- “Relax”
- “Love”
- “Calm down”
- “Ignore AI”
- “Dump the AI shit except for mundane coding”
- “Prioritize system observability”
- “Evolve”
- “Persevere”
- “Boycott AI”
- “Breathe”
- “Panic”
- “Come together”
- “Go outside”
- “Focus on solving real problems, not just chasing AI trends”
- “Recharge”
- “Keep human judgment central to every decision”
- “Stop”
- “Let it happen”
- “Calm the heck down”
- “STUDY”

05 What does the agentic near future hold?

The good and the bad reactions net out to a whole lot of positivity. When asked to reflect, engineers do express worry, but it pales next to their optimism about their roles and personal futures.

What can teams in the present take away from the words of those at the frontier? That successful teams are less worried and less stressed even while facing the same rate of errors. Their success compounds as they use more tools, search more widely for help troubleshooting, and build more of their own developer applications. This may be a look at what some are calling “the great separation” between avid adopters and the rest, as results beget more, automatically, ad infinitum.



Make your AI agents more reliable

Develop AI agents 50% faster with an orchestrator that handles the tricky problems:

- + Maintain state over long periods
- + Ensure state for multi-agent workflows
- + Trigger human interventions
- + Use Temporal in Python, Typescript, Go, and more

Get started for free

[Run it locally →](#)