




2027 REPORT

The AI Exploit Gap

Table of contents

<u>Introduction</u>	3
<u>Summary of key findings</u>	4
<u>Reshaped roles, burned-out security teams: how AI is weakening defenses</u>	5
<u>Attacks are happening, and at an increasing speed</u>	7
<u>Knowing the risks, but missing the fix</u>	11
<u>Prevention is the only solution</u>	15
<u>Leadership recognizes the risks, but lacks urgency</u>	18
<u>Conclusion</u>	21
<u>Methodology</u>	21
<u>About Chainguard</u>	22

INTRODUCTION

The cost of a widening security gap

AI has made nearly every part of building (and attacking) software faster. Teams write and ship code faster. Dependencies accumulate faster. Vulnerabilities get discovered, chained, and exploited faster. The one step that hasn't accelerated to match is fixing those vulnerabilities. Remediation still moves at the speed of human review, and every day that pace holds, attackers gain ground.

That mismatch is “the AI exploit gap” — the widening distance between how fast AI-powered tools can find and chain vulnerabilities and how fast human security teams can respond. Nine out of ten security leaders say AI-accelerated attacks have changed their job more than anything else in the past decade. Over half say burnout is affecting their team's ability to defend against those same threats. The people responsible for closing the gap are being worn down by it.

Security teams also can't close it on their own. When engineering falls behind on patching, the risk moves downstream to security operations and incident response teams, who are left to detect and contain exploitation after it happens. Those teams already face a triple threat: frontier AI models discovering new vulnerabilities faster than ever, threat actors using AI to speed up their attacks, and supply chain disruptions that introduce risks from outside the organization. Closing the gap means engineering and security have to work as partners and share ownership of remediation.

The stakes of getting this wrong are also rising. When AI can find a flaw in minutes, leaving a known vulnerability unpatched becomes harder to defend as a reasonable risk. The EU's [Cyber Resilience Act](#) and [product liability laws](#) make software makers directly liable for faulty products. Carrying thousands of open vulnerabilities in production can now lead to denied insurance claims, shareholder lawsuits, and FTC scrutiny.

Through a survey of 1,300 director-level and above security leaders across nine global markets, we found that many organizations haven't adjusted to this new reality. Even among security leaders who call their organization's security tools “very effective” at stopping AI-accelerated exploits, 77% still say it's likely an attacker will exploit a vulnerability in their organization within the next year. Yet, few share the urgency AI experts have been signaling. And the gap in urgency matters: leaders who feel the most pressure about AI outpacing their teams are doing the most to close vulnerabilities, while those who think they have time to spare are leaving the most exposure open.

The report walks through where that gap is widening, why organizations are still leaving known vulnerabilities unpatched, why teams' tools and processes haven't kept pace, and what's standing between security leaders and board-level urgency. Full [methodology](#) appears at the back of this report.

Summary of key findings

96% of teams feel **security team burnout** has some level of impact on their organization's **ability to defend against AI-accelerated threats**, and **54%** feel the impact is moderate or significant.

91% of respondents agree that AI-accelerated cyberattacks have **reshaped their role** more than anything in the past decade.

76% of respondents say their organization has experienced, or they suspect it has experienced, an **AI-driven attack within the past 12 months**, while **10%** aren't sure or don't have a way to know.

More than three out of four cybersecurity leaders (**77%**) believe it's **likely an attacker will exploit a software vulnerability** at their organization in the next year.

77% of security leaders who **call their tools "very effective"** at stopping AI-accelerated exploits still say **it's likely an attacker will exploit a vulnerability** in their organization within the next year.

93% of security leaders are at least somewhat concerned that their **remediation speed can't keep up with AI-assisted exploits**.

Among leaders who believe AI already outpaces human remediation, **20% report leaving critical vulnerabilities unpatched at least sometimes**; that climbs to **50%** among those who expect AI to outpace remediation within a year, and **64%** among those who think it's more than a year away.

Only **32%** call their **incident response team "very equipped"** for the AI-driven increase in threat volume.

Only **22%** of companies are using **AI-powered security tools "extensively"** to find their own weaknesses before attackers can.


Only **22%** rate their organization's **current security tools "very effective"** at preventing AI-accelerated exploits.

93% of security leaders overall agree their organization's **leadership needs to show more urgency** about AI-accelerated threats (but among CISOs, only **85%** agree).

IMPACTS OF AI-ACCELERATED CYBERATTACKS

Reshaped roles, burned-out security teams: how AI is weakening defenses

The job security leaders have today is not the job they originally signed up for. Nine in ten leaders say AI-accelerated attacks have changed their role more than anything else in the past decade. That shift is landing on people, translating into stress, overwhelm, burnout, and losing the staff organizations need most as threats accelerate. The AI exploit gap is no longer just a technology problem: it's a human problem.



Nine out of ten agree: AI-accelerated cyberattacks have reshaped security roles more than anything in the past decade

91% agree that AI-accelerated cyberattacks have reshaped their role more than anything in the past decade.

Ask a security leader whether AI has changed their job, and the answer is close to unanimous. 91% agree (either “strongly” or “somewhat”) that AI-accelerated cyberattacks have reshaped their role more than anything else in the past decade. Only 1% “strongly disagree” that AI has impacted their role.

Burnout is impacting a majority of organizations’ ability to defend against AI-accelerated threats

64% of security leaders say their concern about AI-accelerated cyberattacks has contributed significantly **to their own sense of burnout** in the past year.

96% of leaders feel security team burnout **has some level of impact** on their organization's ability to defend against AI-accelerated threats, and **54%** feel the impact is moderate or significant.

Despite promises that AI would make security teams' jobs easier, the shift in roles is showing up as strain instead — both personal and organizational. In our study, 64% of security leaders report their concern about AI-accelerated cyberattacks has contributed noticeably (“a moderate amount” or “a great deal”) to their own sense of burnout in the past year. Their feelings aren’t unfounded: in Darktrace's 2026 State of AI Cybersecurity Report, “87% say AI is significantly increasing the workload of their security operations team.”

Our data shows that burnout has serious implications for these companies’ security. Nearly all security leaders (96%) say burnout is having some level of impact on their team's ability to defend against AI-accelerated threats, and 54% agree the impact is moderate or significant.

This burnout may have started before AI sped up vulnerability discovery. While vulnerability remediation typically sits with the engineering or DevOps team as a first defense, security teams often take the blame for open vulnerabilities, or get pulled in to handle detection and response after an exploit. As we’ll cover later, 54% of respondents say their organization knowingly leaves high or critical vulnerabilities unpatched at least sometimes due to capacity and business constraints. Burnout may be the consequence of the security team being held responsible for risk they don’t control.

Impact of security team burnout on organization’s defense against AI-accelerated threats




IMPACTS OF AI-ACCELERATED CYBERATTACKS

Attacks are happening, and at an increasing speed

We've just discussed how the shift to AI-accelerated attacks is impacting the people defending against them. Now let's explore why: the threat itself is growing, and organizations already know it, whether or not they have already experienced an attack.

Three out of four security leaders believe it's likely an attacker will exploit a software vulnerability in their organization within the next year. Three out of four report a confirmed or suspected AI-linked incident in the past 12 months. Roughly half say that within a year, AI will outpace their ability to respond — or already has. Cracks are opening in multiple directions, at a faster rate, and security teams are on the line to keep up. Yet, while they report burnout, these security leaders don't feel the same urgency about AI-accelerated attacks as the industry experts who are sounding the alarm.



Three out of four companies have faced an AI-driven attack

For most organizations, an AI-driven attack is not a hypothetical. In our survey, 76% of security leaders say their organization has likely experienced a security incident involving an AI-discovered or AI-generated exploit in the past 12 months. Of that group, 17% confirmed actual incidents with either attacker tooling, forensic evidence, or vendor/researcher attribution, and another 59% suspected attacks based on signs like an unusually fast time-to-exploit after a patch release or attack sophistication. Another 10% say they are “not sure” whether an attack has happened or not, or they have “no way to know.” Only 14% say “definitely not.”

76% of respondents say their organizations **have experienced, or they suspect they have experienced, an AI-driven attack** within the past 12 months.

10% One in ten respondents say they are **“not sure / no way to know”** if their organization experienced an AI-driven attack within the past 12 months.

Three out of four security leaders see attacks on the horizon

This first-hand experience with attacks is shaping expectations going forward. 77% of security leaders say it's likely an attacker will exploit a vulnerability in their organization within the next year, including 16% who say it's “very likely.” Only 5% call it “not at all likely.”

77% of security leaders say it's likely an attacker will exploit a software vulnerability within their organization in the next year.

77% of security leaders who call their tools “very effective” at stopping AI-accelerated exploits still say it's likely an attacker will exploit a vulnerability in their organization within the next year.

We see some interesting splits. Newer leaders and those at lower revenue tier companies feel more exposed: among security leaders who have been in their roles for more than 10 years, 66% say an attack is likely within a year, whereas among those with a decade or less of experience, 79% say it's likely. 84% of those at companies with under \$250M revenue say an attack is likely within a year, whereas only 72% of those at companies worth over \$1B feel that way.

What's notable is that splitting respondents by title doesn't change the sense of urgency: 77% of C-suite leaders call an attack likely, the same proportion as respondents overall.

Preparedness should, in theory, lower how likely an attack feels. In a surprise turn, our data shows that among leaders who rate their current security tools "very effective" at stopping AI-accelerated exploits, 77% still say an attack is likely — identical to the overall total. In other words, those who have invested time and energy in building their defenses may have become more acutely aware of the threat.

Security leaders aren't as concerned about AI-accelerated attacks as are those closest to AI

It's not about whether an attack will happen, but about whether the people hired to stop it can still keep pace once it does. The world has been warned repeatedly by those closest to AI that AI-powered vulnerability discovery will quickly outpace human teams' ability to remediate. In November 2025, Google Cloud's Cybersecurity Forecast 2026 for the year ahead anticipated threat actors moving from occasional to routine use of AI. In August 2026, more than 100 technology and cybersecurity companies — including Google, Microsoft, CrowdStrike, and AI developers OpenAI and Anthropic — signed an open letter warning that the window to defend against AI-enabled attacks is narrowing as models worldwide become increasingly capable.


That drumbeat of warnings hasn't created a strong sense of urgency among many of the people hired to act on it. We asked security leaders when, if ever, they expect AI-powered vulnerability discovery to routinely outpace human security teams' ability to remediate vulnerabilities. Only 11% say "it already does." Add in the leaders who expect it "within six months" (25%) or "within six to twelve months" (19%), and only about half (54%) expect AI to outpace human remediation within a year. At the far end of the scale, 20% of security leaders think it's three or more years away. So while the alarm is coming from inside the AI industry, most security leaders haven't yet answered it.

Again, we see that newer leaders and those at smaller, lower-revenue companies believe AI-powered vulnerability discovery will outpace their team's ability to remediate sooner.

- Among security leaders who have been in their roles for more than 10 years, 37% expect this to happen within the year, whereas among those with a decade or less of experience, the number is 56%.
- By revenue, 79% of leaders at companies under \$250M expect it within the year, significantly higher than 44% of leaders at companies over \$1B.
- And by company size, 64% of leaders at small organizations (500–999 employees) expect it within the year, compared to 35% at larger companies (5,000+ employees).

When will AI-powered attacks routinely outpace human security teams

Overall C-suite



Looking at the C-suite, we see the story flips. Just 6% of C-suite leaders say AI-powered discovery already outpaces their team's ability to remediate, compared to 11% of security leaders overall, and C-suite respondents are more likely to place that shift further out: 33% expect it within one to two years, versus 26% overall.

Security leaders who are confident their team can patch a zero-day vulnerability within a day are **twice as likely** to believe AI-powered vulnerability discovery **already outpaces** human teams' to remediate, indicating the more-prepared leaders may have better insights as to the uphill battle they are fighting.


Just 6% of C-suite leaders say AI-powered discovery already outpaces their team's ability to remediate, compared to **11%** of security leaders overall.

SPEED IS NOW THE DECIDING FACTOR

Knowing the risks, but missing the fix


Every organization has vulnerabilities waiting to be patched; the question is which ones get fixed and which get shelved for “later.” Our data shows that in many cases, organizations are choosing not to act, even when they know the risk. 54% percent of security leaders say their organization knowingly leaves a high- or critical-severity vulnerability unpatched at least sometimes. And our data shows that while smaller companies are struggling more to patch, large companies are simply choosing not to.

In addition, we’ve noticed a powerful correlation. Two important metrics consistently track in reverse: when security leaders believe AI-powered attacks will routinely outpace human security teams, they leave less vulnerabilities exposed, and vice versa.



Smaller companies are significantly more likely to struggle with patching

Reasons organizations knowingly leave vulnerabilities unpatched



Why are vulnerabilities left unpatched? 86% of respondents cite at least one capacity constraint (difficulty testing patches, not enough staff or resources), and 85% cite at least one business tradeoff (lower priority than other work, too disruptive to operations, revenue-generating work taking precedence). Nearly everyone points to both.

Company revenue is correlated with which reason takes precedence. Companies with under \$250M in revenue are significantly more likely to struggle with capacity constraints, such as difficulty testing the patch before deployment (59% vs. 45% for high-revenue companies). On the other hand, higher-revenue companies (\$1B or more) are more likely to make business trade-offs, listing reasons for leaving vulnerabilities as:

- Revenue-generating product work wins out over patching (46% vs. 26% for lower-revenue companies)
- Vulnerability prioritization criteria haven't been updated for AI-accelerated threats (47% vs. 33% for lower-revenue companies)
- It's too disruptive to business operations (48% vs. 41% for lower-revenue companies)


Over half of companies leave high or critical vulnerabilities unpatched at least sometimes

Before AI-accelerated threats became common, companies often left low- and medium-severity vulnerabilities unpatched. Anthropic's April 2026 research on Claude Mythos Preview shows why that's now more dangerous than ever before: in testing, the model repeatedly chained together several individually minor vulnerabilities — for example, using one bug to defeat a kernel security protection and a separate one to gain memory access — to build a fully functional exploit capable of full system takeover.

In our data, more than half of security leaders say their organization still knowingly leaves vulnerabilities of any severity unpatched at least sometimes: 58% for low- or medium-severity vulnerabilities, and 54% for high- or critical-severity ones. Company revenue is one of the clearest predictors of whether a vulnerability will be patched: companies with \$1B+ in revenue report leaving vulnerabilities unpatched at more than double the rate of companies with under \$250M in revenue.

That gap likely isn't about capability. Large enterprises may be more likely to have formal risk management programs, and many may be choosing to accept the risk, justifying a backlog of open vulnerabilities as a "risk-based" approach. That reasoning is getting harder to defend. The EU's Cyber Resilience Act and its new Product Liability Directive hold software makers directly liable for defective products. In the US, where most cyber litigation relies on negligence law, the question is what a reasonable organization should have done. As AI makes it faster and cheaper to chain vulnerabilities into a working exploit, what counted as due care a few years ago starts to look like negligence today.

Companies that knowingly leave high/critical vulnerabilities unpatched at least sometimes



Leaders who see AI risk as an immediate threat leave their organizations less exposed

We sliced our data many ways and noted that urgency and exposure move in opposite directions: the more a group believes AI-powered vulnerability discovery already outpaces human remediation, the less likely that same group is to report leaving vulnerabilities unpatched. And the markets and company sizes treating the AI exploit gap as a future problem are the same ones with the most exposure sitting in it today.

Among security leaders who believe AI already outpaces human remediation, 20% report leaving critical vulnerabilities unpatched at least sometimes; that climbs to 50% among those who expect AI to outpace remediation within a year, and 64% among those who think it's more than a year away.

Among security leaders who believe AI already outpaces human remediation, **20%** report leaving critical vulnerabilities unpatched at least sometimes; that climbs to **50%** among those who expect AI to outpace remediation within a year, and **64%** among those who think it's more than a year away.


By company revenue

Companies under \$250 million in revenue report both the highest perceived urgency (79% expect AI to outpace human remediation within a year) and the lowest unpatched rate (27%). Companies earning \$1 billion or more report the opposite: the lowest urgency (44%) and the highest unpatched rate (66%).

REMEDICATION CAN'T OUTFRAN THE THREAT

Prevention is the only solution

Security teams can only patch or respond to a vulnerability after it's discovered — prevention is the one lever teams can pull before discovery or an exploit exists. But buying tools that promise to prevent attacks is the easy part. Deploying them, and trusting them, is where organizations are falling behind. Our data shows that the pace of the threat is outrunning confidence in the team, the tools, or both.



Being "very equipped" doesn't mean being fast enough

Speed is the first place the gap shows up. 93% of security leaders report feeling at least somewhat concerned that their remediation speed can't keep up with AI-assisted exploits, with 44% feeling either "extremely concerned" or "very concerned."

While vulnerability management is a frontline defense that typically sits within the engineering and DevOps teams, security leaders are increasingly worried that those frontlines won't hold in the face of AI-discovered vulnerabilities. At the same time, they also lack confidence that their incident response team (IRT), the emergency responders you call in when the front line has broken down, is ready either. The majority of security leaders (66%) call their IRT "somewhat equipped" for the AI-driven increase in threat volume. Only 32% call their IRT "very equipped." And even that smaller, more confident group isn't reassured about outcomes — among the 32% who say their IRT is "very equipped," 78% still say an attack is likely at their organization in the next year.

If a zero-day exploit were discovered today, only **19%** of organizations could patch it within a day. **51%** would require 1-3 days.


93% of security leaders are at least somewhat concerned that their remediation speed can't keep up with AI-assisted exploits.

Among the **32%** of leaders who say their IRT is "very equipped" for AI-driven threat volume, **78%** still say it's likely an attacker will exploit a vulnerability at their organization within the next year.

Rollout and confidence in security tools lags

Nearly every organization has done something in response to AI-accelerated threats, but most of that action stops short of full deployment of new security tools. Over 99% report taking at least one AI-specific action, most commonly reallocating staff (51%) or adopting secure-by-default software components (50%).

Actions taken specifically to address AI-accelerated cyberthreats



Across the board, adoption of AI-powered tools is shallow. Only 22% are using AI-powered security tools "extensively" to find their own weaknesses before attackers can; 55% are still in limited or pilot use. The divide shows up by revenue, too: organizations with under \$250M in revenue are more likely to have adopted AI-powered tools in some form (84%) than those with \$1B+ in revenue (75%).

Even where tools are in place, leaders aren't sold on them. Only 22% rate their organization's current security tools "very effective" at preventing AI-accelerated exploits. And as mentioned earlier, 77% of security leaders who call their tools "very effective" at stopping AI-accelerated exploits still say it's likely an attacker will exploit a vulnerability in their organization within the next year.

Awareness and willingness to invest clearly aren't the problem. Our data shows that organizations aren't executing on their plans and purchasing doesn't mean deploying: just 28% have evaluated, purchased, and rolled out new tools specifically because of AI-accelerated threats in the past 12 months, while 39% have purchased tools they haven't rolled out yet. And 27% of the respondents highly concerned about their remediation speed still haven't purchased or rolled out any new tools — only modestly better than 33% overall.

Only **22%** of companies are using AI-powered security tools "extensively" to find their own weaknesses before attackers can.

Only **22%** rate their organization's current security tools "very effective" at preventing AI-accelerated exploits.


27% of security leaders who are concerned that their organization's remediation speed can't keep pace with AI-assisted exploits still haven't evaluated, purchased, or rolled out any new security tools in response to AI-accelerated threats in the past year.

KEEPING PACE ISN'T THE SAME AS KEEPING UP

Leadership recognizes the risks, but lacks urgency

Budget is not the metric that predicts whether an organization is exposed. Seven in ten security leaders (70%) say securing budget for AI-driven threat defense is either “very” or “somewhat” easy. Great, so problem solved? Not exactly.

Nearly half of security leaders (48%) say their company’s security investment falls short of the level of protection they believe their executives and board of directors expect them to deliver. And 93% of security leaders agree leadership itself needs to show more urgency in protecting the organization from AI-accelerated cyber threats.



Budget is no longer where the AI exploit gap begins

Most security leaders aren't fighting for resources. Seven in ten security leaders (70%) describe the process of securing budget for AI-driven threat defense as easy. For most, the budget conversation itself isn't the blocker.

Company revenue shows a similar split: security leaders at companies with under \$250 million in revenue report an easier process securing budget for AI-accelerated threat defenses (80% easy) than those at companies with \$1 billion or more in revenue (66% easy) — a 14-point gap.

The budget is available, but has it actually kept pace with AI-accelerated threats? 71% of security leaders say their organization's investment in AI-driven threat defense has either kept pace with or stayed ahead of the actual scale of the threat. Just under three in ten (29%) say it has fallen either somewhat or far behind.

Lower-revenue organizations aren't just quicker to secure budget, they're also more likely to say that their budget has kept up. Among companies with \$250M in revenue or less, 79% of security leaders say investment has kept pace with or stayed ahead of the threat, compared to 67% at companies with \$1B or more in revenue.

However, we see that even the leaders most confident in their investment in AI-driven threat defense aren't meaningfully less concerned about an attack happening. Among the security leaders who say their organization's investment has “stayed ahead of” the actual scale of the threat, 77% still say it's likely an attacker will exploit a vulnerability in their organization within the next year.

77%


of security leaders who say their organization's investment in AI-driven threat defense stayed ahead of the actual scale of the threat still say it's likely an attacker will exploit a vulnerability in their organization within the next year.

Leadership awareness isn't the problem. Urgency is.

While most security leaders feel it's easy to ask for security budgets for AI-accelerated threat defense, and most feel they are keeping pace with or ahead of the actual scale of the threats, security leaders' satisfaction with their budget falls short of their executives' and board of directors' expectations. Nearly half of security leaders (48%) say their organization's security investment falls short compared to the level of protection their leadership expects it to deliver.

Company revenue drives a wide split here. Security leaders at companies with under \$250M in revenue are far less likely to say investment falls short of the level of protection leadership expects (30%) than those at \$1B+ revenue companies (55%) — a 25-point gap. The same smaller-revenue group is also far more likely to say investment exceeds what's needed (29%, more than double the 13% overall). Overall, security leaders at smaller-revenue organizations report an easier budget process, investment more likely to have kept pace with the threat, and less of a gap between what's funded and what leadership expects.

Investments in security by company revenue tier



Funding a defense budget and treating the situation as urgent aren't the same thing. Asked directly about urgency more broadly, security leaders are nearly unanimous: 93% agree that their organization's leadership needs to show more urgency to protect against AI-accelerated cyber threats, including 30% who "strongly" agree. On the surface, leadership and the security team appear aligned on what's needed next.

But that consensus isn't evenly held once responses are broken out by title. The C-suite — the group setting the tone on urgency — is the least convinced that more urgency is needed. Just 85% of C-suite respondents agree leadership needs to show more urgency, compared to 93% of respondents overall. C-suite respondents are more than twice as likely to disagree outright. The people closest to defending against AI-accelerated threats see the need for urgency more clearly than the leaders whose job it is to set the pace.

85% of C-suite respondents agree leadership needs to show more urgency, compared to 93% of respondents overall.

FINAL THOUGHTS

Conclusion

From our data, it's clear the AI exploit gap isn't confined to one realm of a security team's job. The gap shows up in how frequently attacks are happening (and will happen), in vulnerabilities organizations know about but haven't patched, in tools bought but not deployed, and in the distance between what security teams say they need and what their boards treat as urgent. And, the gap impacts the people doing the work — burnout that's no longer just personal, but starting to affect how well teams can defend against the threats causing it.

Preparedness, on its own, doesn't seem to change the math. Across the data, leaders who rate their tools highly, their incident response teams as well-equipped, or their defense strategy as keeping pace with the threat all still expect an attack at close to the same rate as everyone else. The leaders who see the AI exploit gap most clearly aren't the least capable ones: they're often the ones paying the closest attention — although even their attention lags the urgency signaled by the AI industry itself.

That exposure is getting harder to explain away. For years, a backlog of unpatched vulnerabilities could be a risk-based decision — a calculated bet that attackers wouldn't find the weak spots before the team did. Now that machines can now find and chain flaws faster than human teams can review them, the odds that a known vulnerability gets exploited keep climbing. And when one does, the record showing the team knew about it and left it open becomes hard to defend to a regulator, an insurer, or a jury.

Closing the gap means treating prevention as the standard, not the aspiration — fixing what's exploitable before it's exploited, rather than racing to catch up after the fact. Security can't do that alone, and engineering shouldn't have to either. The teams that write, build, and ship software are on the front lines of this gap, facing the same AI-driven pressure to move faster. The way forward is a real partnership: shared ownership of remediation, secure defaults that don't slow anyone down, and a common understanding that every vulnerability fixed before production is one no one has to chase after an attack.

Methodology

This survey was commissioned by Chainguard and conducted by Wakefield Research, a third-party research partner, to collect responses from 1,300 cybersecurity leaders globally. Respondents were chosen from a variety of industries, and organization sizes and revenue. All respondents had a minimum seniority of Director, and worked for companies with at least 500 employees (except in Singapore and UAE/Saudi Arabia, where the minimum was 250 employees). Respondents were chosen from nine markets: U.S., U.K., France, Germany, Brazil, Singapore, Mexico, UAE/Kingdom of Saudi Arabia (KSA), and Japan, in September 2026, using an email invitation and an online survey.



Prevention-first cybersecurity

Chainguard is the trusted source for open source. By providing engineers and AI agents with hardened, trusted, and production-ready artifacts, Chainguard helps organizations prevent AI supply chain attacks, increase engineering velocity while reducing toil, and maintain continuous compliance. Customers include Fortune 500 enterprises and global industry leaders, including Anduril, Canva, DocuSign, OpenAI, Public Storage, Snap Inc., and Snowflake. Chainguard is venture-backed by leading investors, including Amplify, IVP, Kleiner Perkins, Lightspeed Venture Partners, Mantis VC, Redpoint Ventures, Sequoia Capital, and Spark Capital.

For more information, visit: www.chainguard.dev

Get in touch with Chainguard to learn more.

