Vanta

The State of Trust Report

THIRD EDITION

Table of contents

The State of Trust

What does the state of trust look like today? According to 2,500 business and IT leaders across the globe, there's good news and bad news.

AI is changing all the rules. AI risks are high (page 5) and, alarmingly, adoption has already outpaced expertise (page 8). That leads to gaps that will keep any business leader up at night.

Everyone is feeling the burn, daily.

Leaders are tasked with reducing risk while also demonstrating trust—and their reputation and growth potential depend on it (page 14). What makes this balance even harder is that risks and time spent on compliance keep rising, but budgets simply aren't budging (page 6).

The worst part? Everyone seems to be stuck in a rut, spending more time *proving*—instead of *improving*—security (page 15).

That's not what trust is all about. It's about outcomes, not optics.

It's time for a shift—from manual "security theater" to continuous, automated trust management focused on real security outcomes.

Vanta's third annual State of Trust report shows that it's possible. Despite its challenges, AI is banishing burnout and removing the manual slog that comes with the compliance and risk management practices of Yesteryear. And agentic AI is unlocking even more possibilities (page 20).

Leaders don't have to inherit old burdens and accept them as the cost of doing business. They can forge a better path forward.

Read on to learn how.

Trust in the age of AI

Risk is higher, and ecosystems are leaking

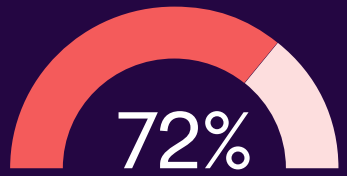

say overall risk is at an all-time high (↑ 17% YoY), and **more than half report a recent vendor breach** (56%).

AI readiness lags adoption —governance is the gap

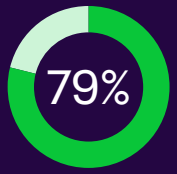

are using or actively planning to use agentic AI

But **65% say their use of agentic AI outpaces their understanding of it**—and only 48% have a framework in place to limit autonomy.

Pressure to prove trust is rising

82%
tie stronger security/ compliance to higher trust

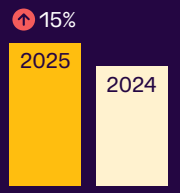

77%
say stakeholders now demand verified proof of compliance

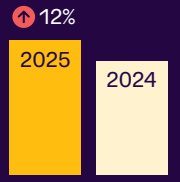

Manual proof work is choking teams

Almost 2/3 say they spend more time proving security rather than improving it.

Teams now spend **12 weeks/year on compliance** (↑ 1) and **9 weeks on vendor reviews** (↑ 2)

Governed automation is the release valve

95%
report gains from AI/automation, with 1 in 2 seeing improved accuracy and faster assessments.

01 Risks are rising, but budgets aren't budging

Today's cybersecurity threats are more frequent, more intelligent, and harder to contain than ever before—with nearly 3 in 4 security leaders saying overall risk has never been higher. It's a huge spike from one year ago, when just 55% said the same.

Continuous threats are the new normal

Why the spike in anxiety among security leaders? They now face non-stop threats. A majority (56%) experience threat activity at least once a week, and nearly 4 in 5 encounter a threat at least once a month. It's a trend that's headed in the wrong direction. In 2024, only 50% reported detection at least once a week.

The frequency of breaches, specifically, is raising alarm bells. Almost half of organizations saw more breach activity in the past year.

With threats on the rise, teams are forced to respond faster and remain in constant firefighting mode. The unrelenting pace of attacks is turning into "threat fatigue" and leading to burnout for security professionals. To thrive (not just survive), security leaders need real-time visibility and faster ways to prioritize threats and avoid backlogs.

Threats are accelerating year-on-year

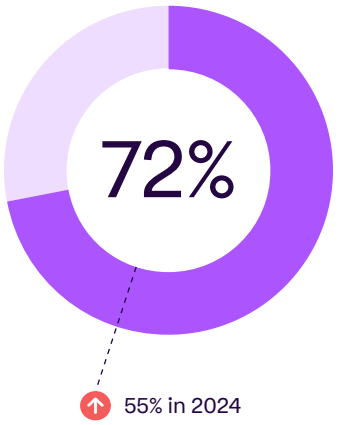

of security decision-makers say risk has never been higher

Organizations face mounting threats

Threat encounter frequency:

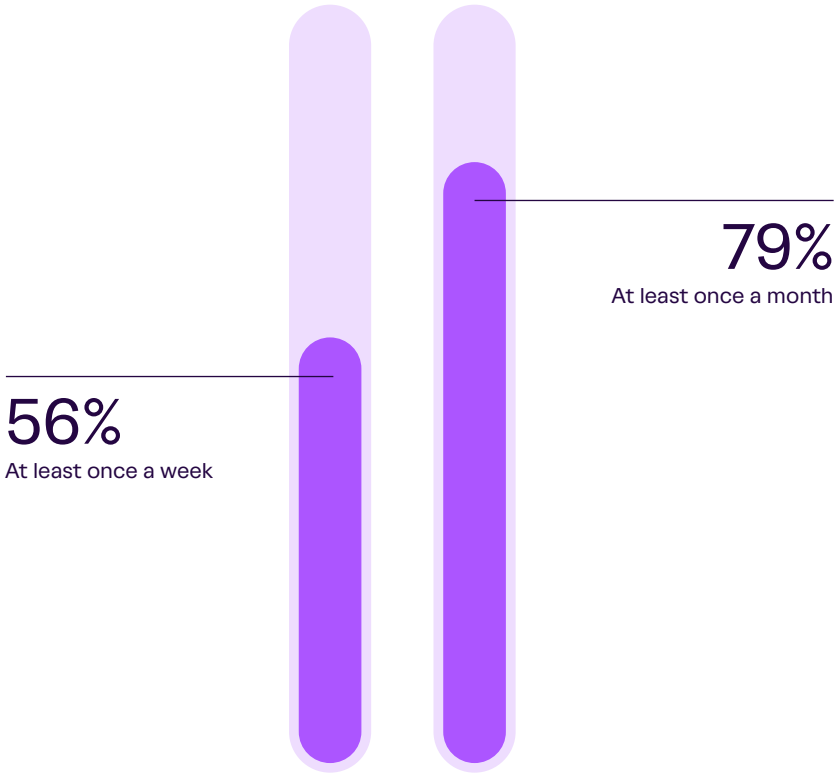

Budgets stuck in neutral

While risks rise, funds remain flat. Leaders say their organizations only spend half of what they should on security—dedicating 10% of IT budgets to security vs a 17% ideal. In 2024, leaders spent 11% vs a 17% ideal, suggesting there's been no progress for strained security teams.

Low IT budgets: the same old story

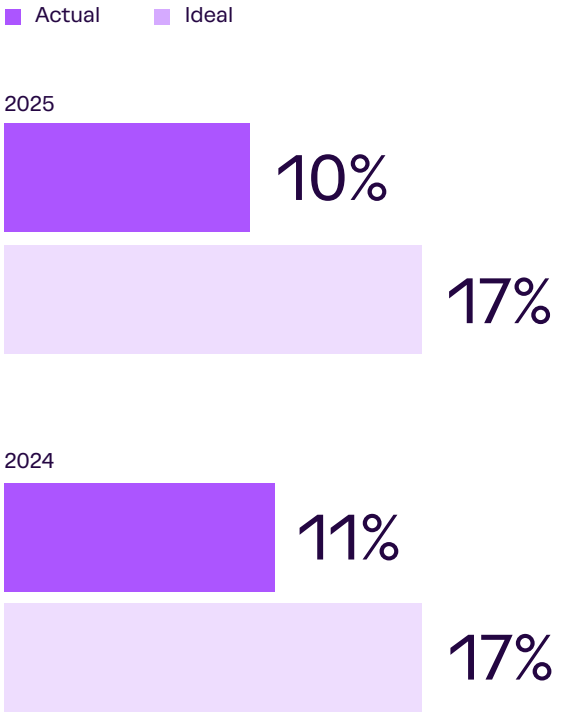

02 The AI readiness gap

Threats aren't just growing in volume—they're changing in nature. Half of all businesses say they've experienced an uptick in AI-generated phishing attacks, AI-powered malware, and AI-driven identity theft or fraud vs last year. Across organizations with over 1,000 employees, the frequency was even higher.

And it's not just volume—it's precision. AI-generated phishing emails are hard to differentiate from legitimate ones, malware morphs mid-execution to avoid detection, and identity theft campaigns can run at a scale previously impossible for hackers.

AI risk has outpaced readiness

For many organizations, it's not just sophistication but speed that's causing problems.

Phishing campaigns that once took weeks to coordinate can now be launched in hours. Fraudulent identities can be spun up, tested, and deployed by algorithms in seconds. This speed puts defenders permanently on the back foot.

The research backs this up. The majority of security teams struggle to manage AI-associated risks, with 3 in 5 (59%) admitting AI-related security threats outpace their expertise. This figure rises to 67% for companies with 1,001-2,000 employees, showing just how big a problem this is for big organizations.

Everyone is experiencing an uptick in AI threats

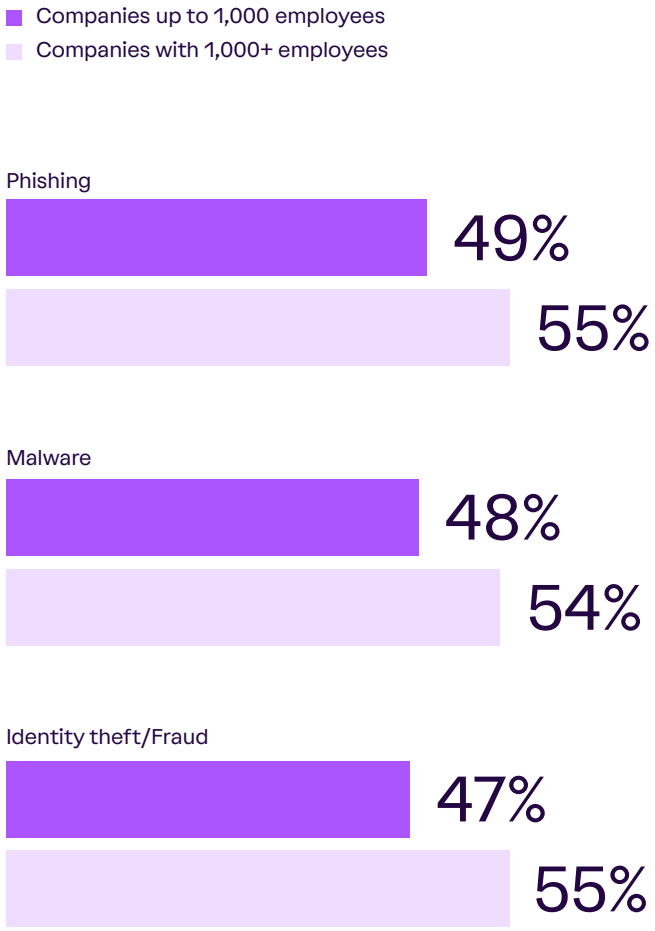

How can teams compensate? Tailored playbooks, adaptive email defenses, model-aware endpoint detection, stronger identity proofing, and rehearsed response procedures designed for synthetic threats.

Yet most organizations don't have these tools in their arsenal. Traditional controls—like signature-based detection, static firewalls, and manual reviews—don't cut it against AI's dynamic attacks. Teams are forced to improvise on the fly, often without the technical training or automated tooling required to respond at the speed AI demands.

Data dilemmas

AI has undoubtedly raised the financial and reputational stakes. And 1 in 3 organizations expect it to lead to increased legal or regulatory exposure.

This is strengthening the case for data usage guidelines. Without them, employees risk feeding sensitive data into public LLMs, unknowingly creating data leaks.

But too many organizations today are being careless with customer data. Naturally, it's making many wary about privacy and trust.

There are some signs of positive momentum. While data protection drags, AI risk assessments and policies are gaining traction this year.

Rules and reviews are gaining traction

44%

have developed and implemented a company AI policy

↑ vs 36% last year

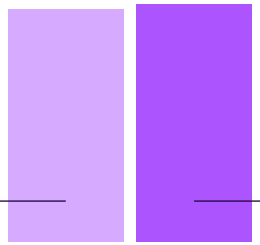

45%

of teams conduct regular AI risk assessments and audits

↑ vs 37% last year

Expertise isn't keeping up

3 in 5 say AI threats outpace expertise (67% in companies with 1,000+ employees)

The governance gap

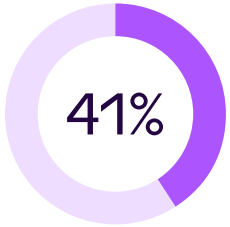

41% Less than half apply strict data minimization

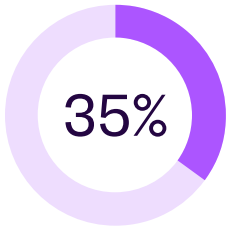

35% Only 35% rely solely on anonymized data to train AI

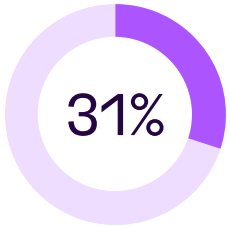

31% Only 31% use customer data to train AI and require opt-in

“To safely implement and manage AI, you have to focus your energy on understanding the risks and how to mitigate those risks using basic controls. It often boils down to re-emphasizing the basics of good security practices, particularly strong access controls. When you have robust security protocols already in place, the adoption of cutting-edge technologies like AI becomes significantly less daunting.”

Mohan Shamachar. Director of Information Security and Compliance

RSI Security

03

Agentic AI adoption is high, but control is low

Today, 8 in 10 organizations are actively using or planning to use agentic AI this year. But a majority (65%) say their current use of agentic AI already outpaces their understanding of it.

Despite this disconnect, agentic AI has moved from experimental to the mainstream. From managing customer queries to making procurement decisions, agents are already embedded in business-critical workflows.

In security, leaders believe it to be most effective for tasks like forensic log analysis and time reconstruction, automated threat correlation and anomaly detection, and generating compliance reports and audit logs.

Agentic AI usage vs understanding

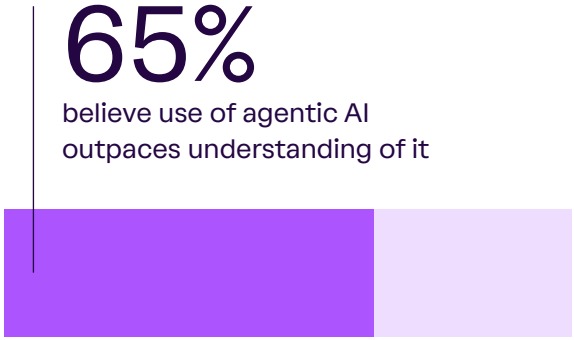

Which tasks would be most effectively handled by an AI agent?

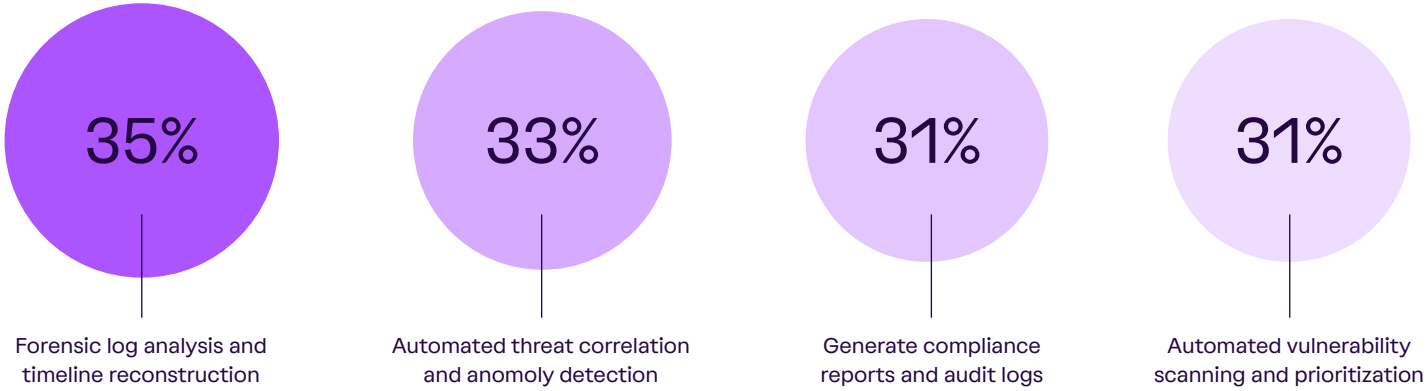

"Vanta isn't just talking about AI—they're building it in ways that genuinely deliver value to my team. As a small company set to double in size, we need tools that help us scale without adding complexity. Vanta's new agentic capabilities are unlocking a level of efficiency and scalability we simply couldn't achieve before."

Braden Pitts, CISO & SVP, Enterprise Technology

The MJ Companies

Organizations are warming up to agentic AI

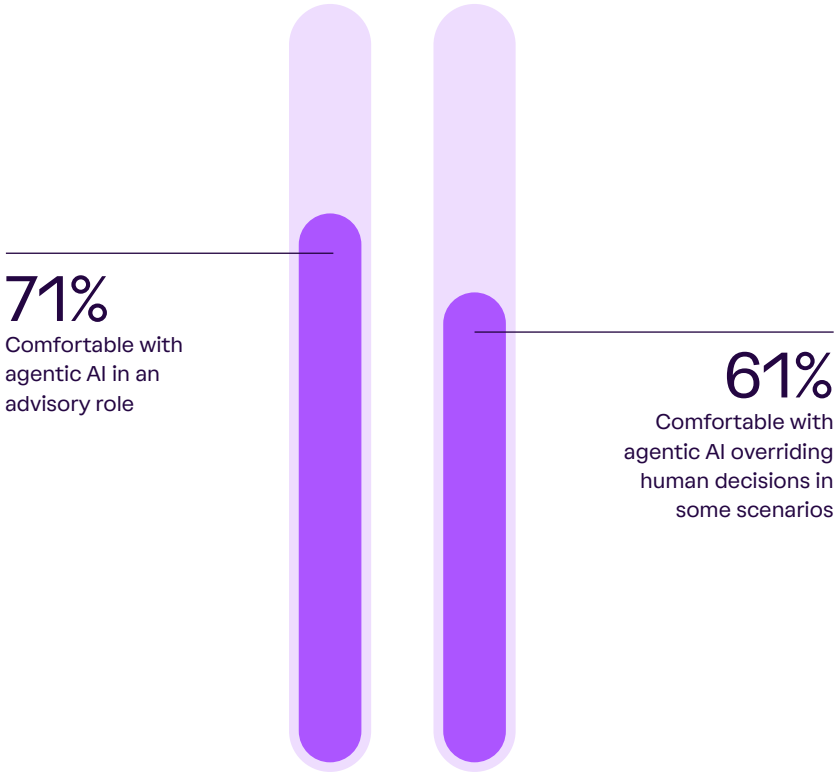

More comfort with autonomy

A whopping 7 in 10 respondents say they would be comfortable with agentic AI providing input on high-level security strategy—with a majority (61%) trusting agentic AI to override a human decision in certain security scenarios.

Of course, this kind of autonomy is appropriate only for pre-approved, time-critical, and reversible controls (e.g., auto-isolate a host during clear ransomware behavior or auto-block a known-bad IP), with immediate human notification and one-click rollback. For irreversible, customer-impacting, or judgment-heavy actions, agents should remain advisory or part of a human-in-the-loop process.

These numbers are a signal of both optimism and fatigue.

On one hand, leaders see agentic AI as a solution to outpace threats and scale security teams. On the other, burnt-out security teams are relieved to offload decisions to machines—even in high-stakes contexts like real-time incident response.

The cultural shift is striking

Not long ago, the idea of machines overriding humans in security settings would have been unthinkable. It wasn't even a decade ago that cloud computing was being called into question.

Now, for the majority, technological reliance is acceptable if the system has pre-approved rules. Is this due to shifting attitudes or improving technology? With 8 in 10 deploying autonomous AI tools, and 71% saying they're comfortable relegating some control to agents, our research shows it's both.

The agentic governance gap

Despite widespread adoption, a minority of companies using agentic AI have a way to control it. Only 48% have developed a framework for granting or limiting autonomy in AI systems. This leaves a huge gap between adoption and control—businesses are embracing autonomy without consistent oversight.

In practice, the consequences of letting agents act with minimal supervision could be scary. Without structured guardrails, decisions made by AI could easily exceed their intended scope. In this reality, a misconfigured autonomous system could lock users out of critical infrastructure, leak sensitive data through poorly trained models, or apply blunt-force controls that disrupt operations. No wonder then that nearly two-thirds of businesses worry agentic AI could erode customer trust.

This is a reputational time bomb for businesses. Customers and regulators are watching how agentic AI is governed, and companies that can demonstrate clear rules, transparent practices, and auditable outcomes will not only protect themselves—they'll have a competitive advantage.

48%

have developed a framework for granting or limiting autonomy in AI systems

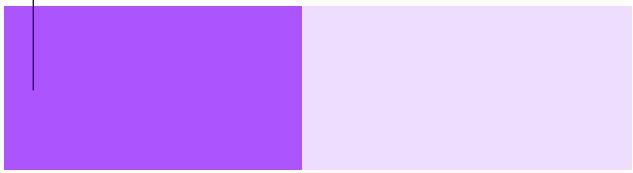

62%

of organizations say agentic AI could erode customer trust

04

Trust is a business imperative

Trust can make or break a business, and security is more important than ever before.

Over 8 in 10 (82%) believe that improving security and compliance directly boosts customer trust—a sharp rise from 67% in 2024. Customers also have higher expectations for trust. This year, 77% of businesses report that stakeholders demand verified proof of compliance, up from just 65% in 2024.

Trust or bust

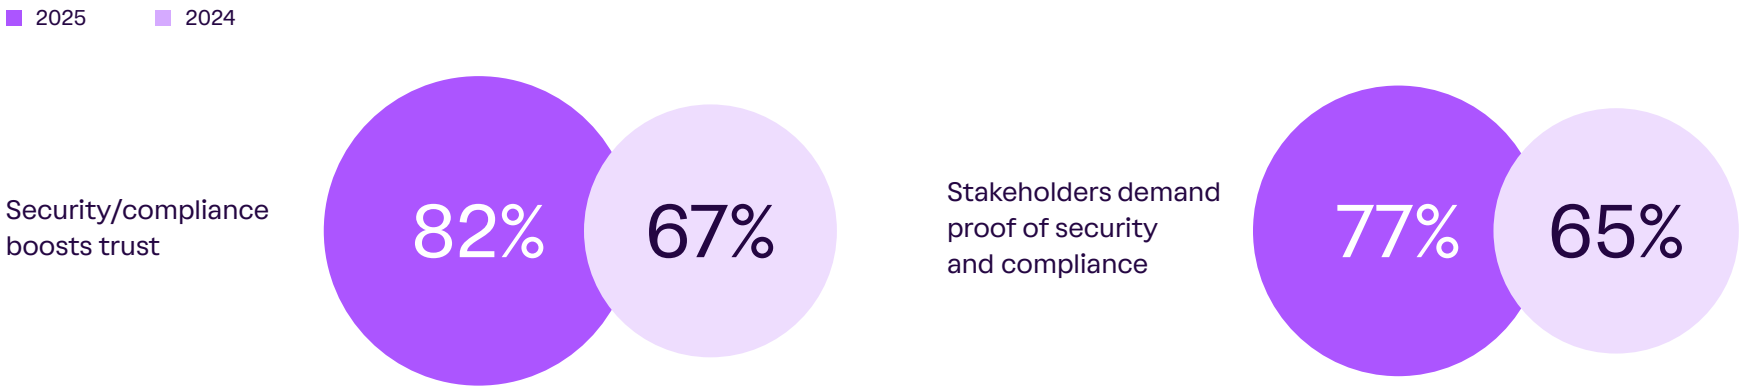

Security is the business advantage

Security is now the visible differentiator in the buying process. It closes deals faster and makes the buying process smoother.

Evidence of strong security can make or break a deal. Procurement teams are quicker to approve vendors with transparent proof of compliance, while hesitation or vague answers slow negotiations to a crawl. For growth-minded businesses, trust and security is no longer a "cost of doing business" but a direct lever for revenue acceleration.

In fact, half of organizations say the biggest value of good security practices is customer trust. While 45% cite improved reputation, 44% mention reduced financial risk, 42% say increased security maturity, and 40% say they meet customer demands for security and regulatory compliance.

The message is clear: customers, partners, and regulators want proof to back up promises.

61%

of respondents say they spend more time proving security rather than improving it

64%

say today's security frameworks feel like security theater

Security theater: Optics over outcomes

But demonstrating that proof is challenging. Many leaders feel the time spent on compliance distracts security professionals from more meaningful work. We've lost the plot, and now almost two-thirds say they spend more time proving security rather than improving it.

It's because teams are buried in manual overhead, spending ~10 hours each week on compliance tasks like policy reviews, evidence collection, and vendor attestations. That adds up to nearly 12 working weeks a year—a whole working week more than 2024.

With this "death by documentation," businesses aren't hardening systems or closing vulnerabilities. They are diverting skilled teams to screenshot-based evidence collection and one-off auditor requests tracked in spreadsheets. The result is frustrated teams, slower innovation, and controls that look good on paper but don't hold up under a real attack.

This is where trust needs to be redefined: not as a one-off certification, but as a living system of record. Companies that can automate processes to provide continuous, verified evidence of compliance will not only win customers faster but will also shift trust from a burden to an advantage.

Buried in manual compliance tasks

“Vanta is a great and reliable platform for people who want to focus on what matters most—security—and not get bogged down in bureaucracy.”

Clementine Markman, Founding Operations Lead

Granola

05

Third-party networks bring more risk and more reviews

Third-party risk makes it harder for organizations to stay secure and demonstrate trust. While two-thirds (67%) feel they have strong or very strong visibility into risk, teams are spending significant amounts of time on vendor security reviews.

And that burden is higher than ever. Today, organizations spend 7 hours per week (an average of 9 working weeks a year) on vendor security reviews and risk assessments—up two whole weeks from the year prior.

Vendor management has become an “always-on” activity. Procurement, renewals, and scope changes trigger another cycle of questionnaires, evidence requests, and follow-ups, creating bottlenecks that strain already thin security resources.

9 working weeks ↑

Teams spend 9 working weeks a year on vendor reviews—up from 7 last year

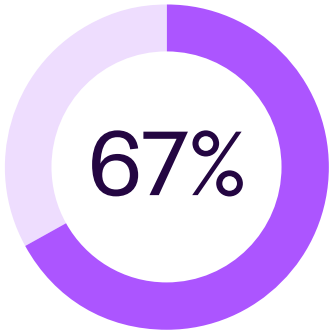

say they have strong visibility into third-party risk—but breaches are more common than ever

Confidence vs lived reality

While 8 in 10 businesses are confident their vendors would inform them of a breach—that doesn't mean breaches are any less common. In fact, a majority (56%) of organizations have had a vendor experience a data breach in the past 6-12 months (up from 48% last year).[1]

Security concerns are also why there's an increase in partners parting ways: nearly 6 in 10 (57%) of businesses have terminated a vendor relationship due to security concerns in the past 6-12 months—up from just 50% last year.[2]

Companies want to believe their vendors are secure—until proof shows them otherwise. When problems come to light, trust disappears fast and vendors are dropped.

For vendors, this means one breach or delayed disclosure can wipe out years of trust. For buyers, it means vendor networks are fragile. One security failure can cascade across the supply chain, forcing sudden exits and major disruption.

The need for continuous assurance

This disconnect highlights a key tension: as third-party ecosystems expand, stronger assurance and continuous monitoring are essential.

Static assessments like annual audits or one-time questionnaires can no longer keep up with the speed of today's threat landscape. Businesses need shared, always-current artifacts, live monitoring of vendor controls, and automated remediation workflows that surface risk before it becomes systemic.

Businesses value vendors —but value trust more

80%
are confident vendors would disclose a breach

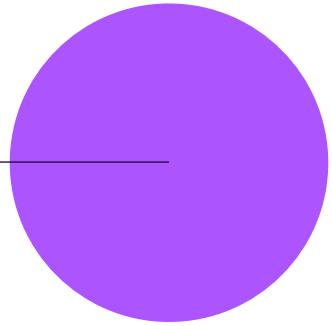

56%
have actually experienced a vendor breach

↑ from 48% in 2024

57%
have terminated a vendor over security concerns

↑ from 50% in 2024

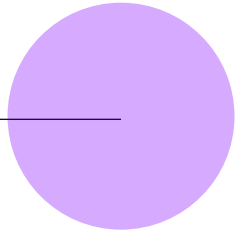

[1] The timeframe for data breaches being reported in 2024 was 'In the past' instead of 'In the past 6-12 months'.
[2] There was no timeframe for a terminated vendor relationship due to security concerns in 2024.

“Managing vendors from a security perspective used to be a lengthy process. We were sending out questionnaires, gathering artifacts, going back and forth, but now, all of this is automated.”

Marciano Kruithof, VP of InfoSec

Bynder

06 The promise of automation and AI

AI and automation are easing the pressure on resource-strapped teams, and almost all businesses believe AI and automation have improved security team effectiveness.

Time for strategy—or a perception gap?

Among IT decision-makers, 53% say that AI and automation have increased time for more strategic, high-level work. Among business decision-makers, that figure drops to 43%.

This difference reveals a perception gap. IT leaders, who directly see reclaimed hours in their workflows—work like incident response, access reviews, and log analysis—are quick to credit automation with freeing them for strategy. Business leaders, further removed from day-to-day operations, often struggle to see those same gains translate into clear business outcomes.

Closing this gap requires showcasing automation's wins in terms executives care about: faster deal cycles, reduced downtime, cleaner audits, and tangible risk reduction.

95%

of businesses believe AI and automation have improved security team effectiveness

Undeniable efficiency

A whopping 8 in 10 organizations are increasing AI usage in security programs to potentially counter rising threat complexity. And it's not hard to see why—AI is supercharging their day-to-day work in several key ways: 51% report faster risk assessments, 50% see improved accuracy, and 36% cite streamlined compliance.

For teams buried under alerts, vendor questionnaires, and evidence collection, automation has become more than a convenience. It's survival. By removing repetitive manual work, AI allows all teams to operate at enterprise scale.

The impact is not just operational but emotional. Over 3 in 4 (76%) respondents say AI is banishing burnout. Long cited as one of the top drivers of security turnover, when staff can hand off the "grunt work" to automation, they can refocus on higher-value, more rewarding tasks.

From stopgap to superpower

The organizations leading in automation use reclaimed capacity to reshape their approach to security. Free from endless compliance checklists, they're able to invest time back into strategic, high-value work—which can look like anything from red-team simulations, to hardening controls and innovating with new AI-enabled defense models.

The takeaway: automation can no longer be framed as a tactical tool for efficiency. It is a strategic enabler that shifts security from reactive firefighting to proactive trust-building, transforming the function from a cost center into a source of competitive advantage.

How organizations are benefiting from AI

“As a rapidly growing security program, Vanta has saved me personally up to 12 hours per week, which lets me put that time towards developing other high priority security objectives that are mission critical for us as a business.”

Mandy Matthew, Lead Security Risk Program Manager

Duolingo

Conclusion: Trust defines survival

Security and trust have become impossible to ignore. They are business-defining.

But rising threats, rapid AI adoption, and the weight of compliance have created an environment where manual, static approaches don't cut it anymore.

The real kicker? Trust is central to growth, but most organizations don't have the budget to back their ambitions with action. Meaning too much time spent on optics over outcomes.

But there's a better path forward. Trust can—and must—be dynamic, measurable, and continuously proven. With tools built for today's reality, businesses can flip the script: evolving security from defensive "theater" to a strategic advantage.

The opportunity is there. Are you ready to take it?

It's your move. 3 ways to turn these insights into action.

01

Leverage automation and AI

AI is supercharging security teams and banishing burnout. Look for tools that use automation and AI agents to eliminate manual work like evidence collection and policy updates—so you can focus on building a strong, scalable GRC program and secure your business.

02

Continuously monitor your vendors

TPRM is getting harder and more time-consuming. Static, point-in-time reviews leave you vulnerable. Continuous monitoring is the key to staying ahead and safeguarding your business.

03

Proactively demonstrate your security posture

Good security drives customer trust. Don't let customers and prospects have any doubts about your program—showcase your security posture through a public trust center for instant and accurate assurance.

Methodology

In July 2025, quantitative research conducted by Sapio Research was commissioned by Vanta to understand the challenges and opportunities businesses are facing when it comes to security and trust management. Vanta and Sapio Research co-designed the questionnaire and surveyed the behaviors and attitudes of 2,500 business and IT leaders across the U.S., UK, and Australia. Tracking data from the 2024 State of Trust Report has also been included, sample sizes in 2024 were 1,000 in the UK and U.S. and 500 in Australia.

Please note, in 2025 the sample also included Supervisors/ Junior Managers, where in previous years they had been excluded. In some instances, question wording or answer options may have changed slightly from previous years. Updates are made to ensure the questions remain relevant and clear. Full details are available on request.

About Vanta

Vanta is the leading AI-powered trust management platform that helps businesses earn and prove trust. Companies including Atlassian, Duolingo, Icelandair, Ramp, and Synthesia rely on Vanta to build, maintain, and demonstrate their trust, all in a way that is real-time and transparent.

Vanta