



The Generative AI Tipping Point

ExtraHop surveyed 1200 IT and security leaders from around the world to understand their plans for securing and governing the use of generative AI tools inside their organizations. Their responses are concerning.



Table of Contents

Executive Summary	3
Key Takeaways	4
Employee Use of AI is High	5
Bans Don't Work, But Remain Common	8
Leaders May Be Overconfident in Their Ability to Secure Generative AI	12
Concerns About Accuracy and Data Exposure	14
Organizations Plan to Invest in AI Security Measures, with Some Exceptions	17
Leaders Welcome Government Guidance	19
Regional Trends	21
United Kingdom	22
United States	22
France	22
Germany	22
Singapore	23
Australia	23
How to Get the Most Out of Generative AI	24
Survey Methodology	27



Executive Summary

Technology and security leaders are once again standing at an inflection point. OpenAI's ChatGPT was released to the public in November 2022, and within four days of launching had more than 1 million users. In contrast to previous waves of technology innovation, like cloud computing, which typically see a slow adoption rate, generative AI and large language models (LLMs) have exploded in popularity. So it was not surprising when our survey found that 73% of respondents said employees in their organization used a generative AI tool or LLM sometimes or frequently. That number will likely only increase as economic and competitive pressures drive organizations to use AI tools in new ways.

What was surprising? An overwhelming majority of respondents (nearly 82%) say they're confident that their organization's current security stack can protect against threats from generative AI tools, yet 74% are planning to invest in generative AI security measures this year. Hopefully those investments don't come too late.

Though employee adoption is high, leaders have their concerns about the technology. OpenAI CEO Sam Altman, in an interview with the [Times of India](#), said his biggest fear is "the hypothetical idea that we already have done something really bad by launching ChatGPT."



These tools have tremendous power, but that only underscores how important it is for the creators and users of this technology to understand its risks.

Raja Mukerji

ExtraHop co-founder and chief scientist

The top concern among survey respondents is receiving inaccurate or nonsensical responses, followed by exposure of personally identifiable information (PII) and the compliance violations that exposure could bring. Apprehension about biases in the tools ranked fourth. Nearly one third of respondents' organizations were so concerned about the risks of generative AI that they banned these tools outright, and half invested in technology that allows them to monitor usage.

Another theme to emerge from the survey findings is cognitive dissonance. Although one in three organizations has banned generative AI usage, only 5% of respondents report that their employees never use them. Presumably, if the bans were effective, these two numbers should be closer, if not equal.

Another notable disconnect: respondents' confidence in their ability to defend against AI threats was high (nearly 82%), despite 50% of respondents not having any technology in place for monitoring employee use of these tools and despite only 42% and 46% offering user training and governance policies, respectively. With just over a third of respondents saying they were highly confident in their security and a similar number saying their organization had banned the use of AI, it raises the question whether bans have led to an inflated sense of confidence.

Generative AI is a mercurial and rapidly growing technology. As such, it's hard to predict how it may change in even six months. In the time since ExtraHop commissioned this survey, OpenAI released an enterprise version of ChatGPT, which purportedly does not leverage user submissions to train its model, alleviating many organizations' concerns over IP loss and exposure.

It's an exciting time for AI and its proponents, but our survey data indicates that organizations have much catch up work to do to ensure their implementations are secure and risks are adequately mitigated.

Key Takeaways

1

Generative AI is here to stay.

Nearly three quarters of global respondents report frequent or occasional use, and 74% are planning to invest in generative AI protections or security measures in 2023.

2

Organizational policies, governance, and training lag behind employee adoption.

While 73% of respondents say employees use these tools with some regularity, only 46% have policies in place governing acceptable use and only 42% train users on safe use of these tools.

3

The data suggests that bans are ineffective.

Almost a third (32%) of respondents say their organization has banned the use of AI tools, but only 5% say employees never use them.

4

IT leaders have their concerns about this technology.

Top concerns include inaccurate or nonsensical responses, exposure of PII, compliance violations, and biases.

5

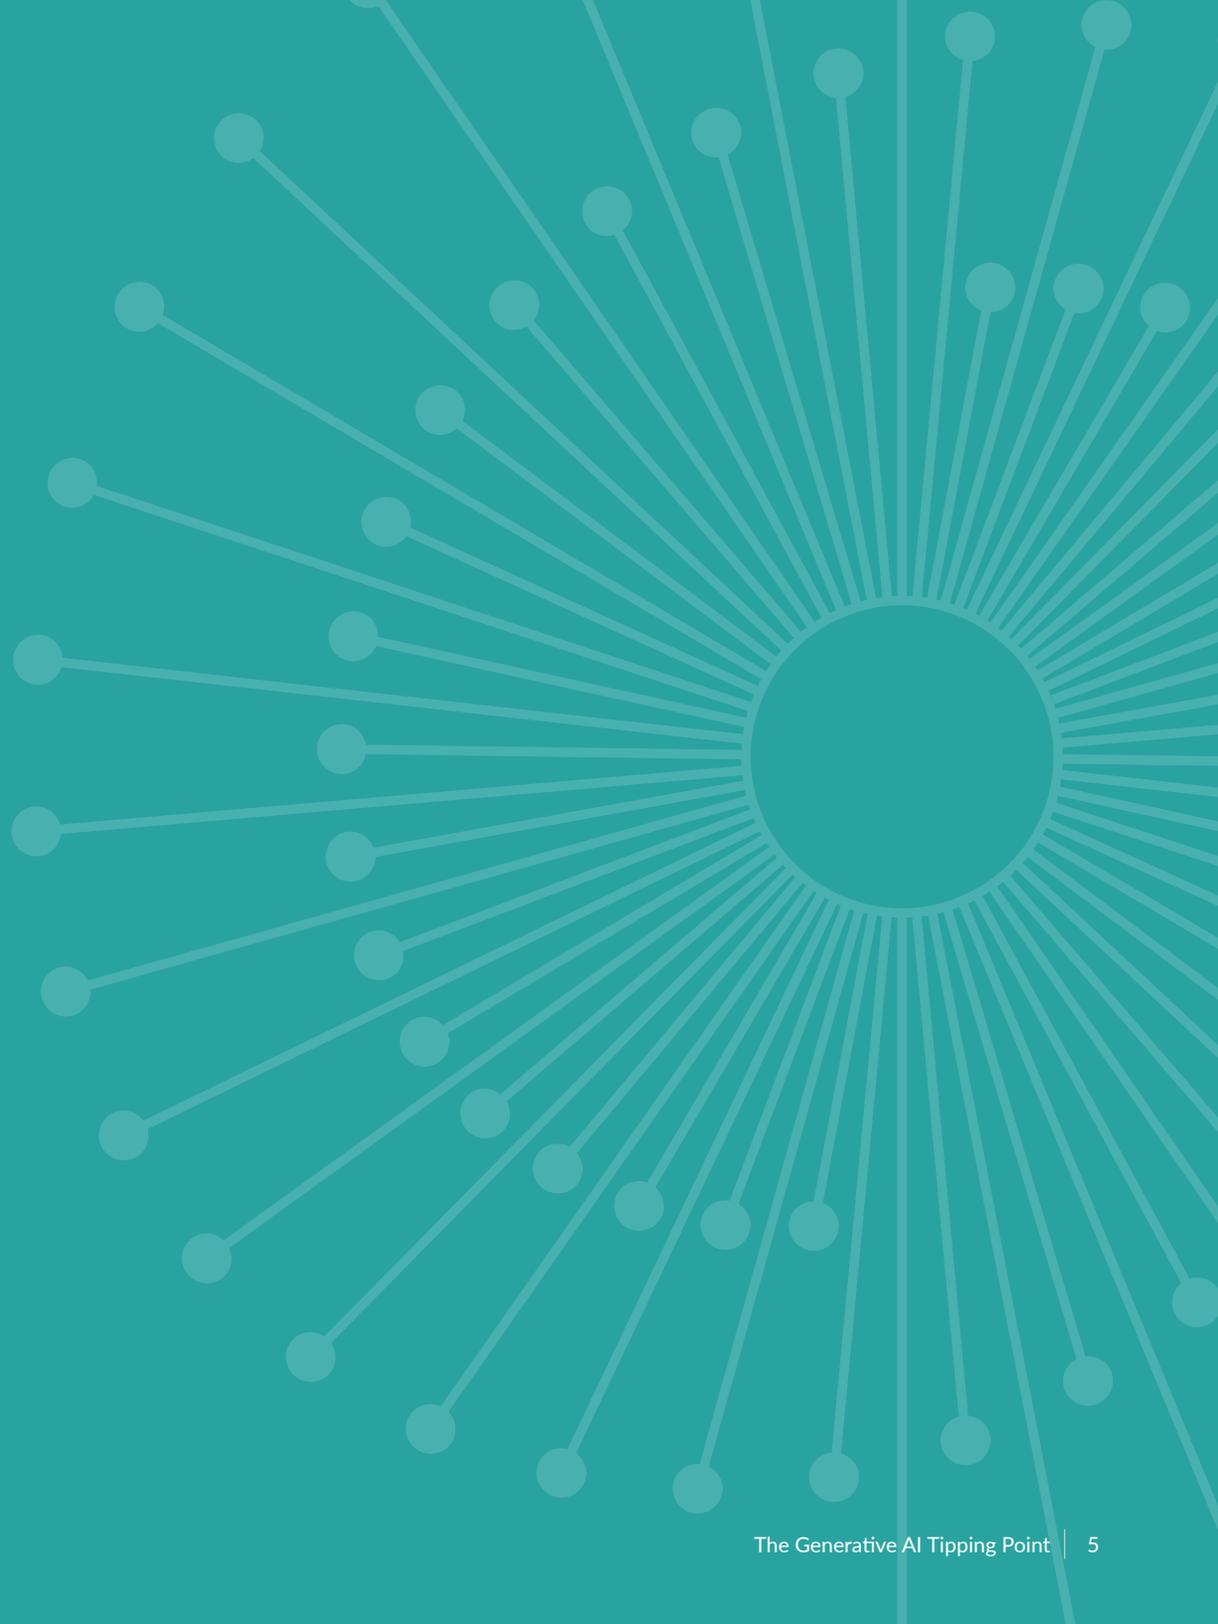
Debate continues on how best to regulate AI.

An overwhelming majority (90%) of respondents say they want the government involved in some way: 60% favor mandatory regulations while 30% support government standards that businesses can adopt at their discretion. Notably, support for government regulation is inversely correlated with age.

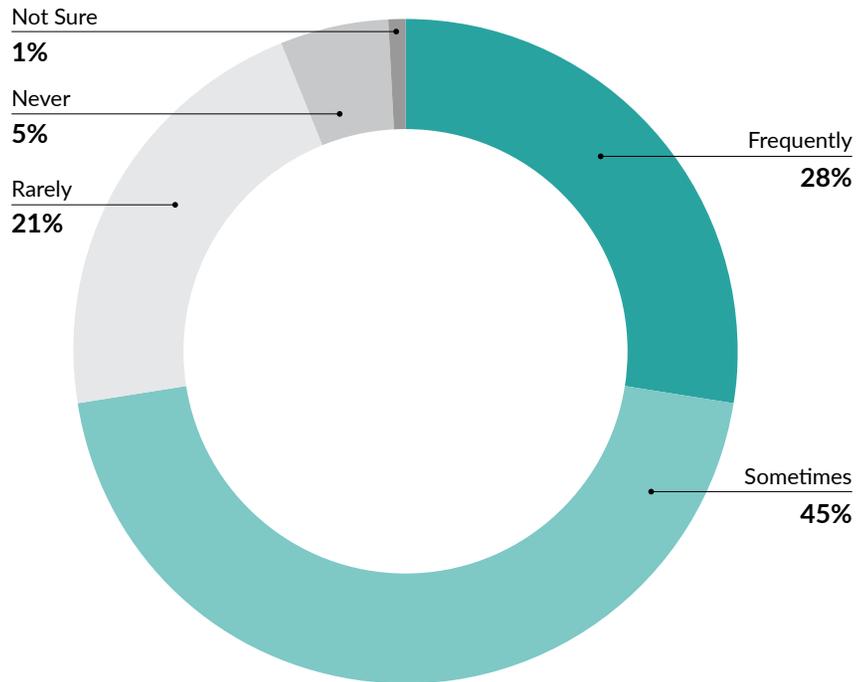
73%

of respondents said employees in their organization used a generative AI tool.

Employee Use of AI is High



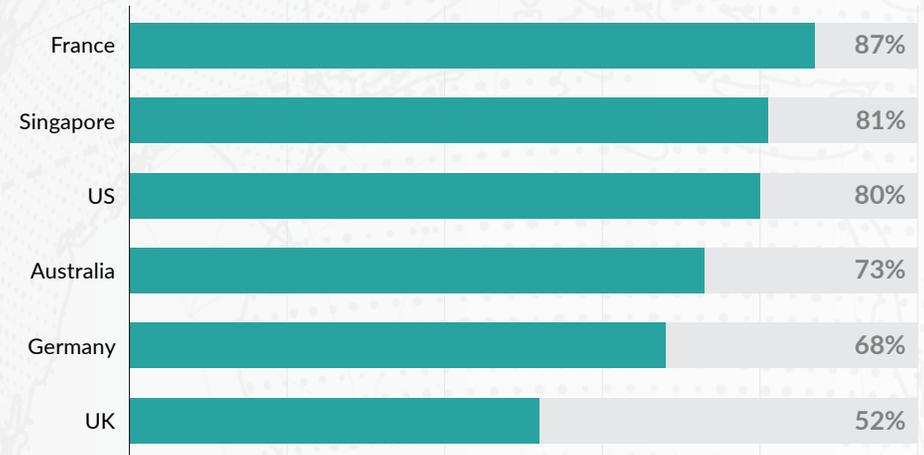
Percent of Respondents Who Say Employees Use Generative AI for Work



Nearly three fourths (73%) of respondents say that employees sometimes or frequently use generative AI or LLMs.

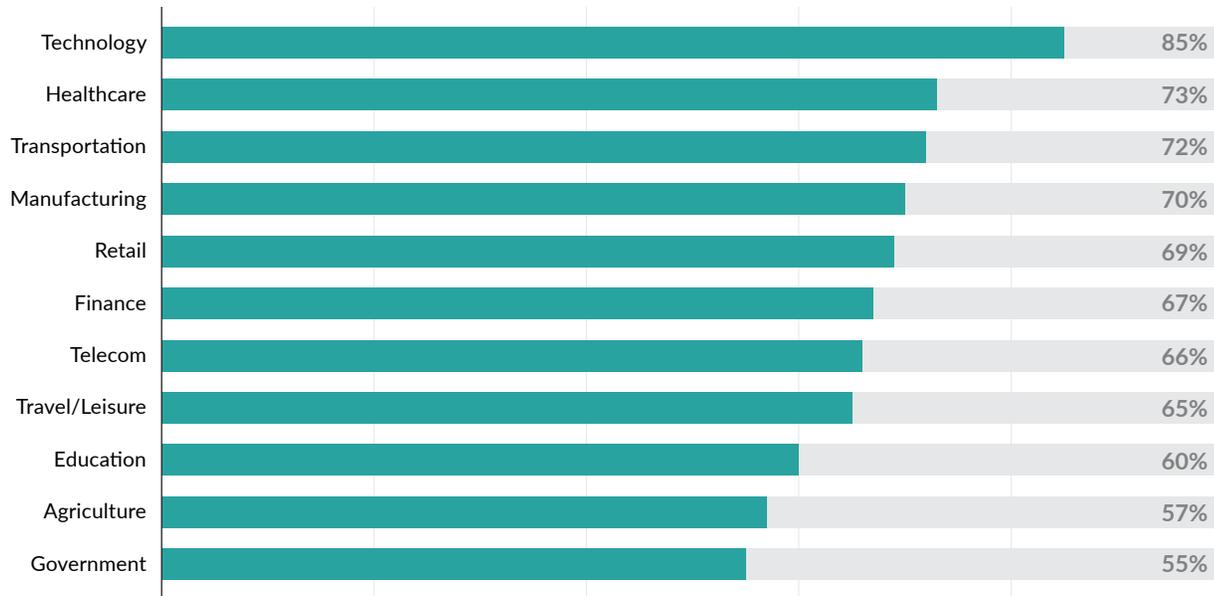
Breaking down AI adoption by country, the top three (in order) are France (87%), Singapore (81%), and the US (80%). Conversely, adoption is lowest in the UK, where nearly half of respondents report employees rarely (35%) or never (11%) use AI tools.

Employee Generative AI Use by Country



Organization size seems to have little effect on the results. Although there was some variation by industry, the majority of organizations are using AI tools with some regularity across the board.

Employee Generative AI Use by Industry

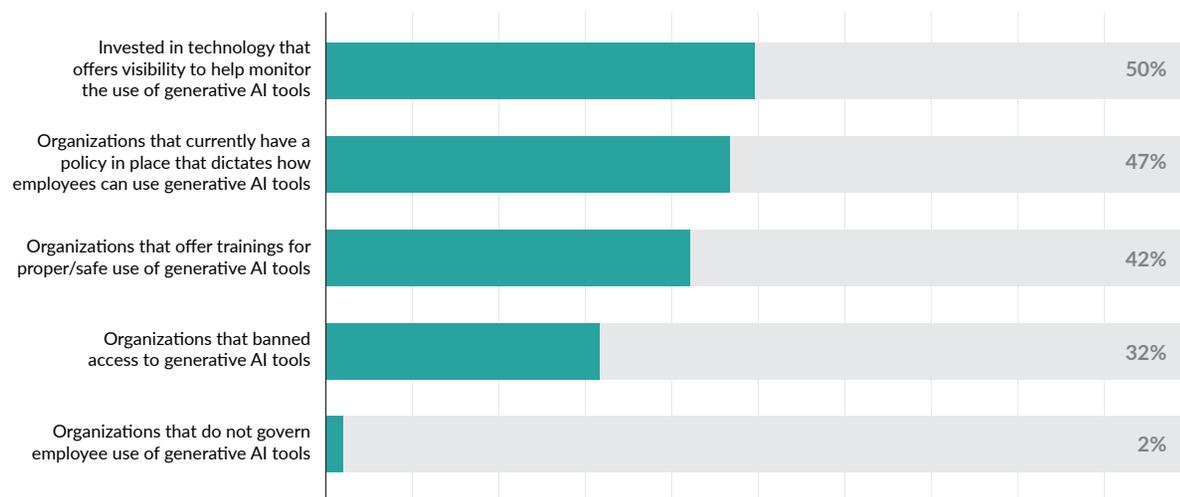


Unsurprisingly, the tech industry leads the way, with 85% of respondents saying employees frequently or sometimes use generative AI tools or LLMs, despite Apple and Verizon announcing bans in May 2023 and others like Accenture, Samsung, and Amazon restricting their use.



Bans Don't Work, But Remain Common

Steps Organizations Are Taking to Govern Employee Use of Generative AI Tools



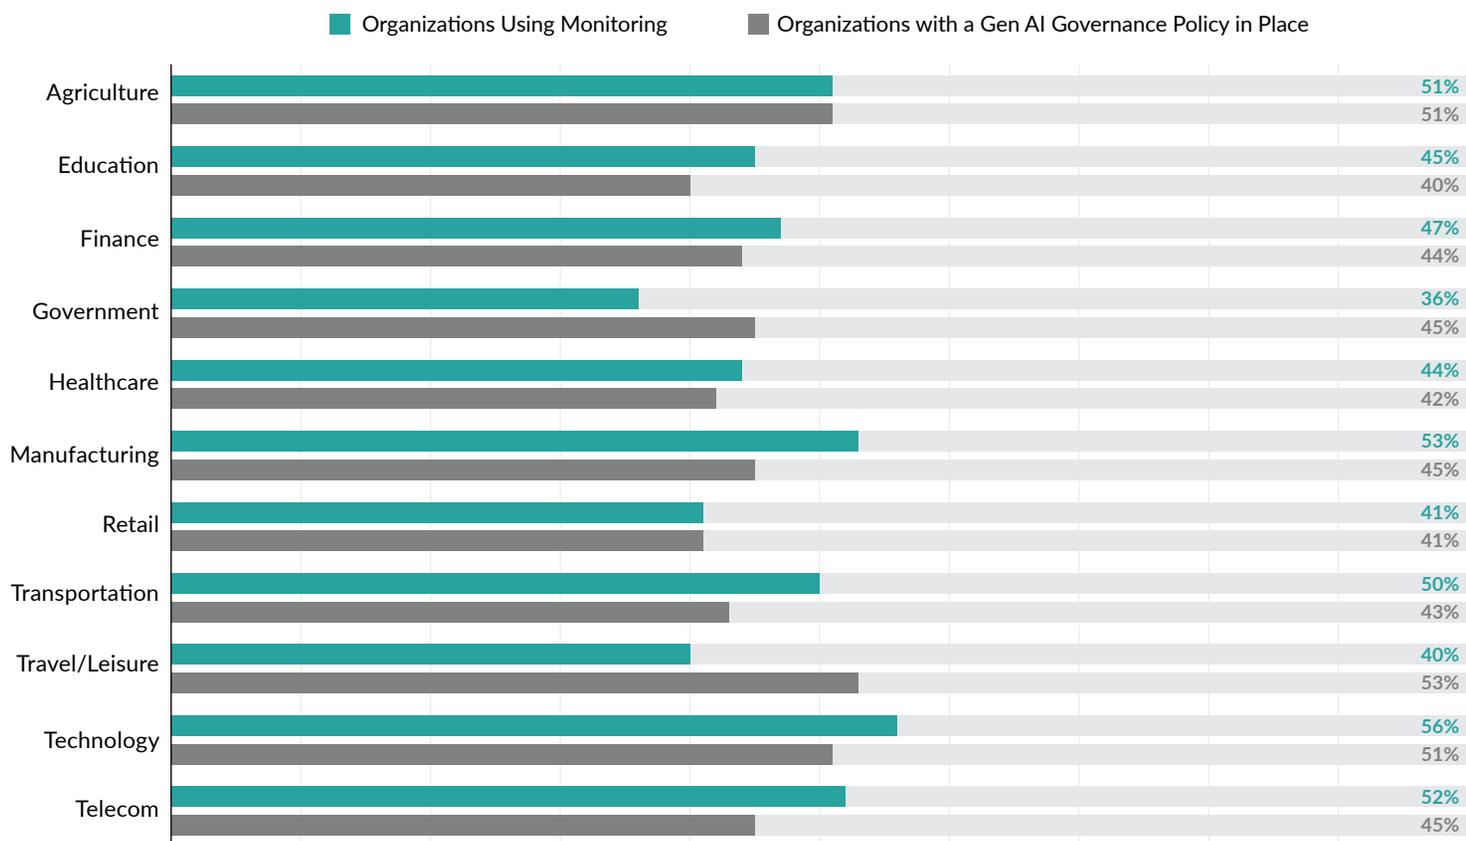
With only 2% of global respondents indicating that their organization does nothing to govern employee use of generative AI tools, it's apparent that most organizations are taking this technology seriously and trying to find ways to maximize its benefits while mitigating its risks. However, it's also clear these efforts aren't keeping pace with adoption rates, and the effectiveness of some of their actions—like bans—may be questionable. Nearly a third of respondents indicate that their organization has banned the use of generative AI, yet only 5% say employees never use AI or LLMs at work. Prohibition rarely has the desired effect, and that seems to hold true for AI.

In what appears to be a trend throughout the survey, the French and the technology sector lead the way in setting the foundation for secure and effective use of generative AI. In France, 58% of leaders say their organization has invested in tools designed to give them visibility into employee use of generative AI, 60% have a policy governing AI use, and 59% offer training around proper use, the highest rate for each. The US (56%) and the UK (50%) are the only other countries where at least half of organizations have invested in visibility tools.



Prohibition rarely has the desired effect, and that seems to hold true for AI.

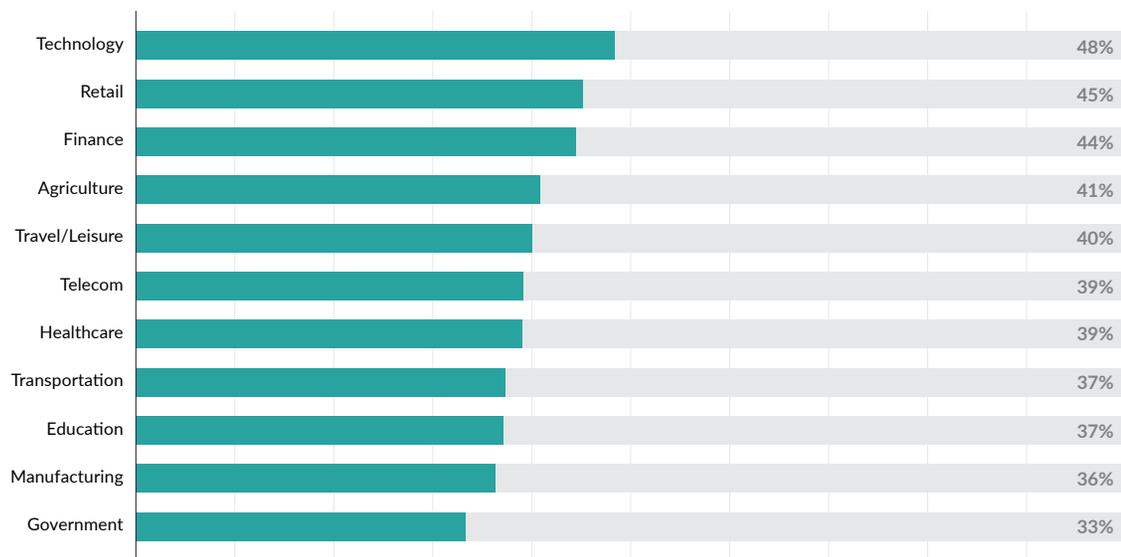
Percent of Organizations That Have Implemented AI Monitoring Technology or Governance Policies, by Industry



It's interesting to note that France also has the highest rate of AI prohibition, with 45% of respondents saying their organization has banned its use. This may suggest that overconfidence or cognitive dissonance is also most prevalent in France, as only 4% said employees never use AI.

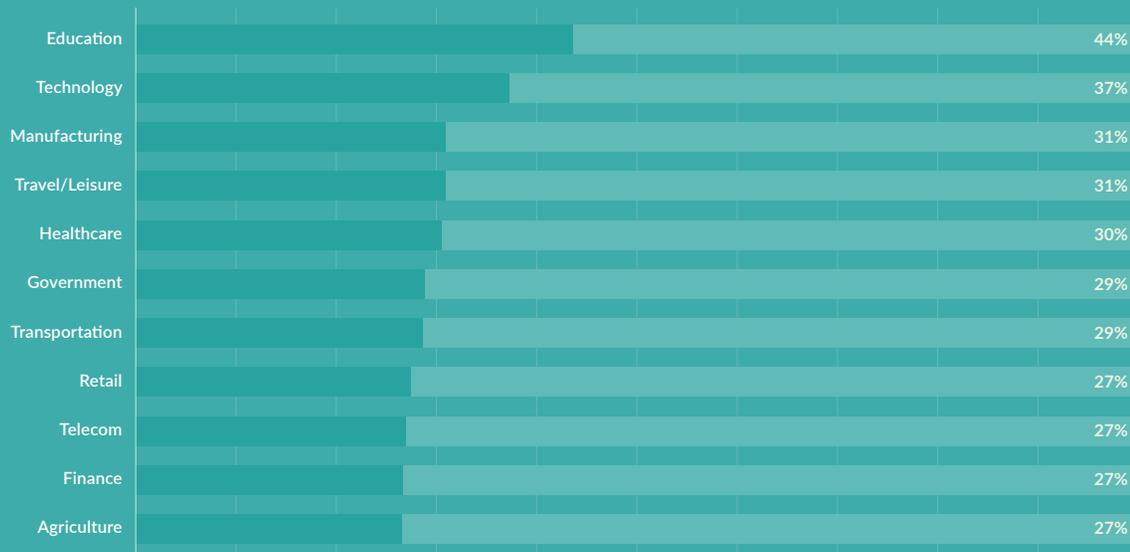
The tech sector is most likely to have visibility tools (56%) and offer training (48%) and ties for third most likely to have a policy governing AI use (51%) with the agriculture sector. At 53%, government organizations are most likely to have a policy in place, followed by travel and leisure. The education sector is most likely to ban AI use entirely, presumably over concerns about plagiarism, ethics, and bias, with 44% saying their organization has done so.

Percent of Organizations that Provide Employees with Acceptable Use Training for Gen AI



Additional findings also appear to support a larger trend of cognitive dissonance. Similar numbers of respondents indicate that they are highly confident in their ability to defend against gen AI threats (36%) and that they have banned the use of AI tools (32%), so it's possible their confidence is tied to the bans their organizations have implemented. But at the same time, just under half of leaders say they have invested in technology that helps them monitor the use of these tools. Thus, at least half of organizations are flying blind, with no way to monitor compliance with policies like bans, yet 82% claim they are at least somewhat confident they can protect their organizations.

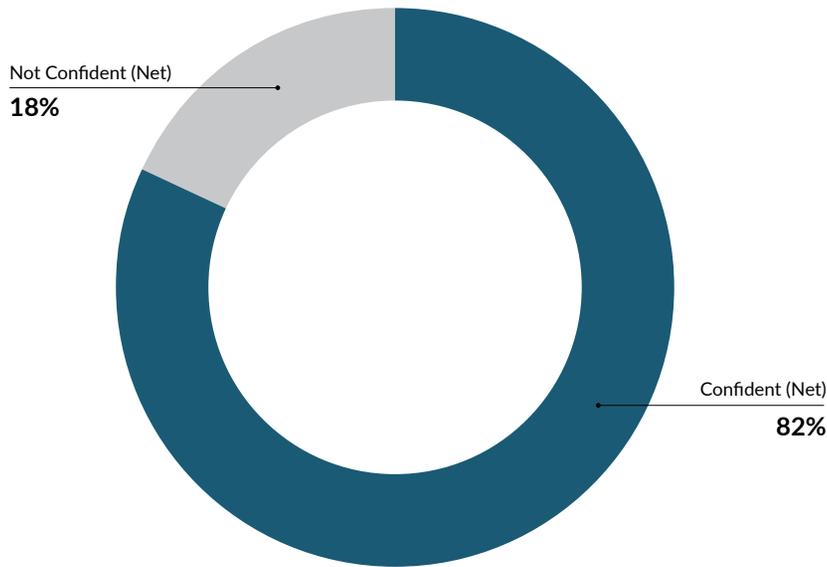
Generative AI Bans, by Industry



At least half of organizations are flying blind, with no way to monitor compliance with policies like bans.

Leaders May Be Overconfident in Their Ability to Secure Generative AI

Percent of Respondents Confident Their Organization's Security Stack Can Protect Against Generative AI Threats



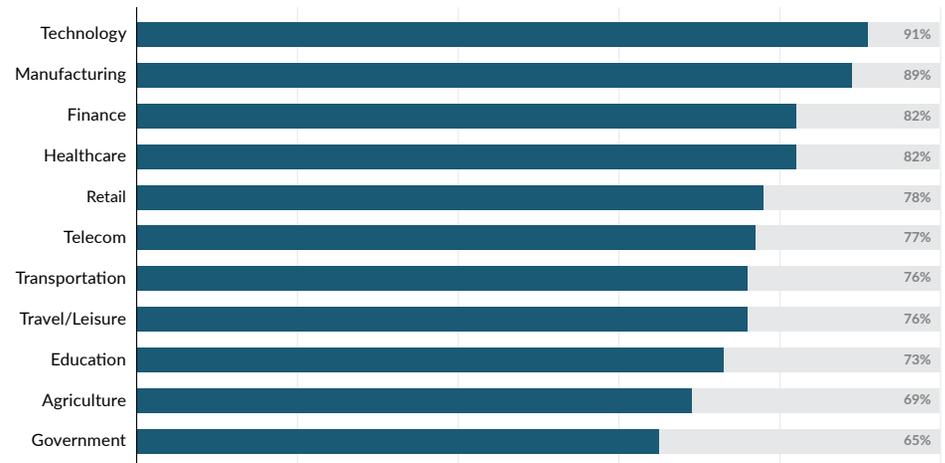
Overall, IT decision makers are confident in their ability to protect against AI threats, with 82% of respondents indicating they are very or somewhat confident.

These results, however, may indicate overconfidence or a misunderstanding of the risks associated with generative AI and LLMs. When respondents were asked how, if at all, their organization governed AI usage, **nearly one third say their organization has banned the use of generative AI tools**, a similar proportion to those who are very confident in their ability to protect against AI threats (36%). But with only 5% of respondents indicating that their employees never use generative AI at work, the numbers don't add up. In other words, if bans were effective, we would expect to see higher numbers of respondents saying employees never use these tools. Even if organizations ban access to these tools, employees can still use them on their personal devices. **And once sensitive information has been submitted, there's no getting it back.**



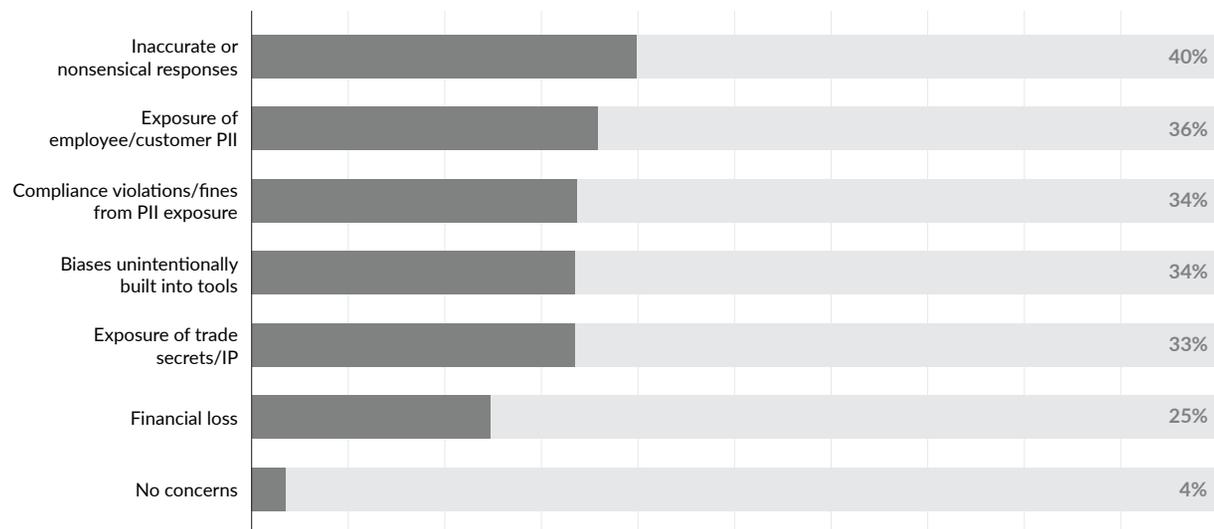
Once sensitive information has been submitted, there's no getting it back.

Percent of Respondents Somewhat or Very Confident in Their Current Security Stack, by Industry



Concerns About Accuracy and Data Exposure

Greatest Concerns About Generative AI Tools



Looking at the global results, the top concern for respondents is getting inaccurate or nonsensical responses, followed closely by exposure of employee or customer personally identifiable information (PII). Less than one quarter of respondents express concern about financial loss.

It's interesting to note where this trend doesn't hold true.

Singapore and Australia are the only countries where the accuracy or coherence of responses is not the top concern. Instead, the most common concern for both countries is exposure of employee or customer PII. For Singaporeans, the second concern is compliance violations and fines,

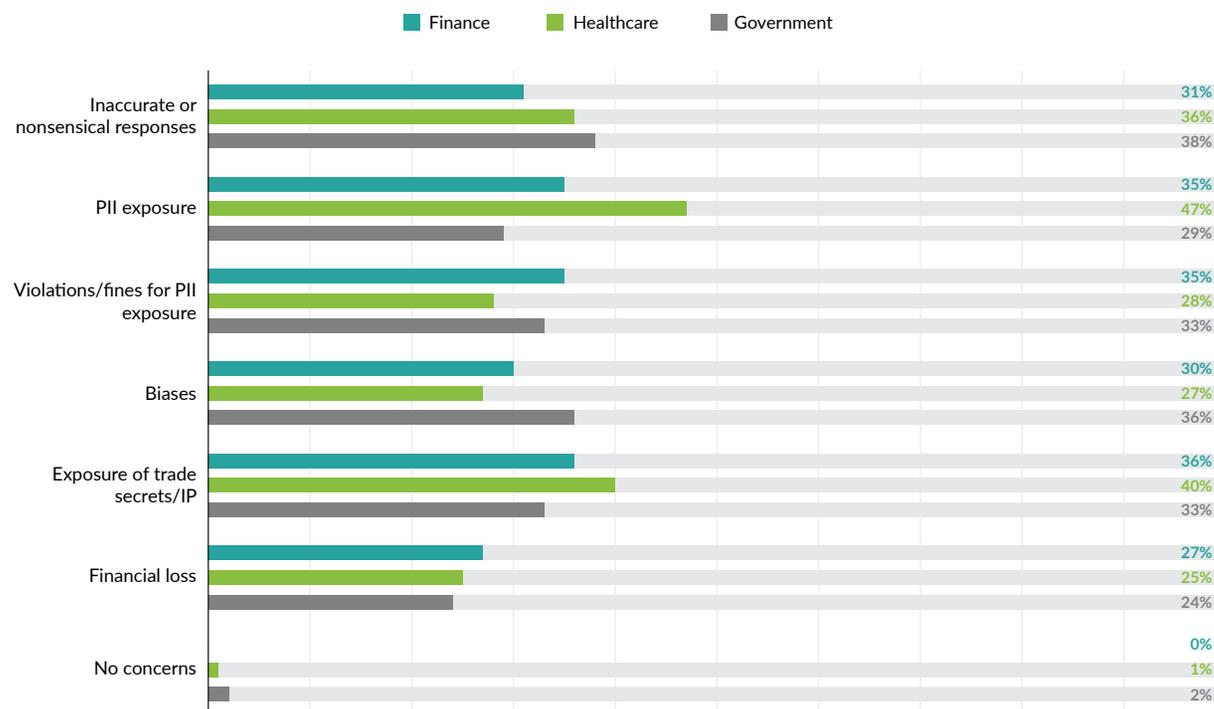
whereas Australians are next most concerned about biases built into these tools. In fact, Australians express more concern about bias than any other region (39%). Though French concerns about bias closely approached this figure (38%), it received the third fewest responses there, ahead of only financial loss and "no concerns."

Similarly, the 55+ age demographic is the only one where accuracy or coherence isn't a top two concern. Once again, the top concern is exposure of employee or customer PII, with just over half of respondents expressing concern about this risk. Of those aged 25-34, the top concern is exposure of trade secrets or IP.

From an industry perspective, finance and healthcare also buck the trend. The financial sector is most concerned about exposure of trade secrets or IP, exposure of employee or customer PII, and compliance violations or fines stemming from PII exposure, with approximately 35% of respondents selecting each of these. For healthcare leaders, the top concern is exposure of PII (47%), followed by exposure of trade secrets or IP (40%). Both finance and healthcare are highly regulated industries subject to serious consequences for exposure of PII. Curiously, concerns about fines or compliance violations from PII exposure didn't match concerns about exposure of this information in the healthcare sector, with only 28% of respondents selecting this response.

While users of these tools can avoid inaccurate or incoherent responses [through better prompting](#), preventing data exposure isn't as simple as many IT and security decision makers may think. Some are turning to their next-generation firewalls (NGFW) to block the IP addresses of these services so that employees can't access them, but NGFWs don't flow nicely into security investigation workflows without a lot of work. **Moreover, blocklists almost always play catch up to end users, and end users can typically find a workaround. Finally, organizations that block the traffic to these domains miss out on the opportunity for monitoring—a mandate for federal agencies in the US—and for auditing compliance with internal use policies.**

Concerns About Generative AI in Finance, Healthcare, Government

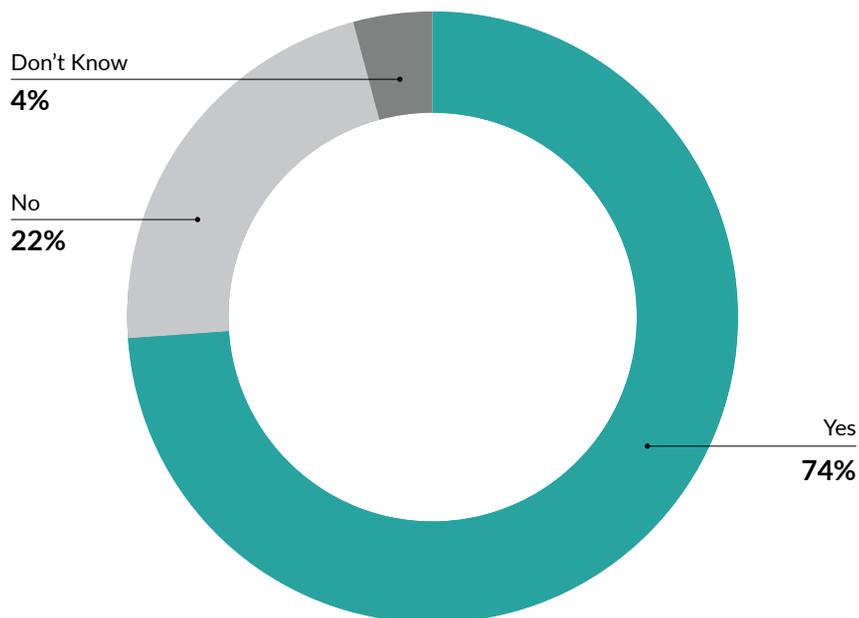


Preventing data exposure isn't as simple as many IT and security decision makers may think.

Organizations Plan to Invest in AI Security Measures, with Some Exceptions



Percent of Respondents Planning to Invest in Generative AI Security Measures in 2023



It seems that greater awareness and use of these tools drives investment in security measures.

Across every region, age demographic, industry, company size, and level of seniority, the response is the same: a resounding yes to the question “is your organization planning to invest in generative AI security measures in 2023?”

There are only two notable exceptions—the United Kingdom and the government sector. Exactly half of government respondents say their organization plans to invest in AI security measures. Meanwhile, only 49% of UK respondents say the same, with 43% indicating there are no plans to invest this year.

The numbers from the UK seem baffling, but adoption is also the lowest here, with only 52% of respondents saying that employees sometimes or frequently use AI tools. A recent Deloitte study also found that only 52% of people in the UK had heard of generative AI, and far fewer (8%) reported using it at work.¹

Adoption is similarly low in the government sector: only 55% of respondents say employees in their organizations frequently or sometimes use AI tools. Government organizations are among the slowest adopters of new technologies, so this could partially explain the low rate of investment in security. American federal agencies in particular, however, should strongly reconsider delaying these investments, as mandates like [Executive Order 13960](#) become more common.

It seems that greater awareness and use of these tools drives investment in security measures, as the countries that lead the way in adoption (France, Singapore, and the US) also overwhelmingly plan to invest in security (96%, 81%, and 84%, respectively).

1. 2023 Digital Consumer Trends. Deloitte. <https://www2.deloitte.com/uk/en/pages/press-releases/articles/more-than-four-million-people-in-the-uk-have-used-generative-ai-for-work-deloitte.html>

Leaders Welcome Government Guidance

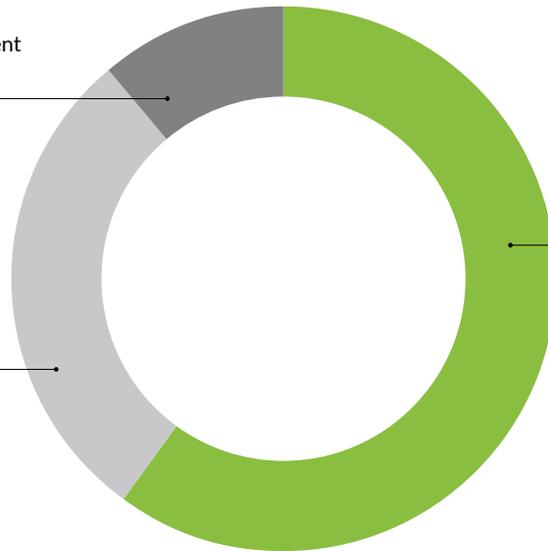
Opinions on Government Involvement in AI Regulations

Businesses should develop their own set of regulations around generative AI tools; the government should not be involved

11%

The government should set clear standards but ultimately leave it to individual businesses to adopt at their discretion

29%



The government should set clear regulations that businesses must follow to protect them from the risks of generative AI tools

60%

60%

of respondents say they believe the government should set clear regulations that businesses must follow.

Across the globe, 60% of respondents say they believe the government should set clear regulations that businesses must follow. French respondents are most in favor of government regulation, with 80% agreeing. The UK is the only region where this is not the majority opinion, as well as the region most likely to say the government should not be involved, with 18% of respondents selecting this answer.

Support for government involvement has a high inverse correlation with age. Nearly 70% of respondents aged 25-34 agree that

the government should set clear, mandatory regulations, but that figure trends downwards to 42% for the 55+ demographic. Similarly, support for optional standards grows from 21% among 25-34 year old respondents to 41% among those over 55 and the proportion of those who say the government shouldn't be involved nearly doubles, from 9% (ages 25-34) to 17% (55+).

Broken down by industry, the tech sector, surprisingly, is most in favor of mandatory government regulation, with 75% agreeing. Typically resistant to regulation, tech companies

may be seeking reassurance that the tools they've invested heavily in won't cause undue risk to their or their customers' operations. The education sector, meanwhile, displays the highest support for optional regulation, with 52% choosing this answer. Those working in transportation (18%) or government (17%) are most likely to say the government shouldn't be involved.

It's clear that leaders want at least some form of guidance from the government when it comes to AI, with 89% of respondents in favor of either mandatory or optional regulations.

Regional Trends



Respondents hailed from the United Kingdom, United States, France, Germany, Singapore, and Australia. Analyzing the data from each country individually yields some interesting findings.



United Kingdom

Britain lags behind the other countries when it comes to AI adoption, with higher numbers of respondents saying their employees rarely (35%) or never (11%) use generative AI or LLMs than any other country. On a similar note, UK respondents expressed the lowest confidence, with 43% not at all or not particularly confident in their ability to protect against AI threats. They also lag behind on investment plans, with 43% of respondents saying their organization has no plans to invest in AI protection in 2023. Additionally, Britons are the most likely to say the government should not be involved in regulating AI, with 18% taking this stance.



France

Of the surveyed countries, France is in the vanguard when it comes to AI. Adoption here is highest (85% use AI frequently or sometimes), as is confidence in security (98%), and plans to invest in security (96%). French organizations have also taken the most action to govern AI tools internally—58% have invested in monitoring technology, 60% currently have a policy dictating how employees may use AI, 59% offer training around proper AI use, and 45% have banned AI outright. Mandatory government regulation also sees its highest support in France, with 80% of respondents in favor.



United States

While not the leader in adoption, confidence, or plans to invest, US organizations do land on the podium for each, behind France and Singapore in adoption and confidence and ranking second in plans to invest. American respondents also seem to be the least worried about AI, with 5% indicating they have no concerns. This confidence seems to have some basis, as American organizations trail only France when it comes to investments in monitoring and AI training.



Germany

German responses largely hew to the average, with a few exceptions. Here, respondents show the highest support for optional government standards (36%). Nearly a third of respondents say they have no plans to invest in AI security this year, and another third said their organization had banned AI use entirely, the second-highest rate for each response.



Singapore

Singapore is another generative AI leader. Adoption is second only to France (81%), as is confidence in AI protections (90%). Singaporean organizations slip to third, behind the US, when it comes to plans to invest in AI protections this year (81%). Support for mandatory government regulations is also high, with 62% in favor. The rest of the numbers seem to indicate a similarly high level of overconfidence or cognitive dissonance, however. Only 40% of respondents say their organization has invested in visibility and monitoring tools, only 39% have internal policies in place, and only 38% offer training.



Australia

AI adoption in Australia matches the global average at 73%, outpacing Germany and the UK. Confidence in protective capabilities also sits close to average, at 79%. Australian concerns, however, are a bit unique compared to other countries'. The top concern here is exposure of employee or customer PII (43%), and more Australians express concerns about biases built into AI tools than any other country (39%). Australians also show the least concern for compliance violations, which is somewhat counterintuitive as they are more concerned about exposure of PII than any other country besides France. One fourth of Australian respondents say their organization has no plans to invest in AI security measures this year, which should be cause for some concern, as Australian organizations have also invested the least in visibility (40%) and offer the least training (34%).

How to Get the Most Out of Generative AI



Generative AI tools and LLMs offer myriad benefits, but they aren't without their risks. Here are 8 helpful tips to get the most out of your implementation:



Establish an internal generative AI task force.

ExtraHop believes that generative AI has the potential to bring enormous productivity gains to organizations, but we also recognize the technology is not without its risks and issues. Therefore, it's worth establishing a cross-functional task force with representatives from IT, security, HR, legal, risk management, compliance, and other functions to explore use cases for the technology inside your organization; to evaluate the pros, cons, and security of different generative AI and LLM tools; and to source training for employees.



Create policies governing safe and effective use of AI.

Policies should include what data can and cannot be shared with public generative AI tools, the circumstances under which AI tools may be used, and how use of these tools should be disclosed to customers. Begin exploring and implementing different technologies that can help your organization monitor use and enforce and audit compliance with internal generative AI policies.



Experiment early and often, but only make moves you can reverse.

AI is still in its early days, so it's important to experiment with different use cases and tools so you can identify what works best for your organization. Your organization's generative AI task force can recommend low-risk use cases where you can start. Engage your organization's legal team to parse the terms and conditions of different generative AI services so that you know what you're getting into, how your data will be used, and how it will be protected. Ask employees to come forward with tools they're using so that you can get some governance around their experiments, and implement technology that will give you visibility into organizational use of generative AI. It's critical to make sure your organization's generative AI pilots won't cause any lasting damage to data or your organization's reputation. After all, once data is in a public generative AI tool, there's no getting it out.



Keep track of where and how your organization is using AI.

With the rapid proliferation of generative AI tools, it's easy to lose track of where and how employees are using them. Some organizations are now facing requirements to track their AI usage, like US federal agencies that must conduct annual inventories and continuously monitor AI tools for safety and effectiveness under [Executive Order 13960](#). Investing in monitoring technology, like [Reveal\(x\)](#), can help you [keep track of both approved and rogue usage of AI tools](#).



Consider building in house if data privacy is a concern.

If your planned use cases involve sensitive data, consider building your own tools or paying for an enterprise license that will give you more control over what happens to user submissions.



Provide training for employees.

Many people don't fully understand how generative AI tools work or the risks associated with them. For instance, many LLMs produce "hallucinations" (seemingly coherent, but nonsensical answers) and inaccurate or biased responses, and some may retain user prompts for training, which could lead to the inadvertent leakage of sensitive or proprietary information. Employees must be made aware of these risks and their potential consequences, especially for organizations in highly regulated industries, and given clear rules of the road for using these tools safely.



Implement other technology and security controls besides blocking.

Many technology and security leaders think they have an easy fix for generative AI: they'll just use their NGFWs to block access to generative AI IP addresses. But blocklists are almost always playing catch up to end users, and end users can typically find a workaround. Moreover, NGFWs don't flow nicely into security investigation workflows without a lot of work in the event of a leak. Finally, organizations that block the traffic to generative AI domains miss out on the opportunity for monitoring—a mandate for federal agencies in the US—and for auditing compliance with internal use policies.



If your organization must prohibit use of generative AI tools, consider a temporary ban and be sure you can enforce and audit compliance with this policy.

Some organizations are banning generative AI because they view it as the simplest way to address the risks. But as our survey findings show, bans are rarely foolproof. Moreover, organizations that ban the use of these tools miss out on the potentially enormous productivity benefits they could realize. A temporary ban, however, may be necessary for some organizations as they navigate the policies and technologies they may need to implement to maximize the benefits of generative AI while minimizing the risks.

Additional Resources on Generative AI from ExtraHop

-  [The Basics of Generative AI](#)
-  [A Harvard “Masterclass” on Artificial Intelligence](#)
-  [Continuous Compromise: Saving AI from Itself](#)
-  [Detect Data Leaks from OpenAI ChatGPT with Reveal\(x\)](#)
-  [Executive Brief | AI Executive Order & EH Gen AI Capability](#)

Survey Methodology

In order to better understand the security challenges organizations face when it comes to employee use of generative AI and large language model (LLM) tools, like ChatGPT and Google Bard, ExtraHop worked with Censuswide in early Fall 2023 to conduct a survey of IT and security decision makers from around the world. Censuswide selected 1200 respondents at the director level or above who worked at organizations with greater than 1000 employees and who influence their organization’s security and IT decisions.

ABOUT EXTRAHOP NETWORKS

ExtraHop is the cybersecurity partner enterprises trust to reveal the unknown and unmask the attack. The ExtraHop Reveal(x) 360 platform is the only network detection and response solution that delivers the 360-degree visibility needed to uncover the cybertruth. When organizations have full network transparency with ExtraHop, they can see more, know more, and stop more cyberattacks.

Learn more at www.extrahop.com