SECOND ANNUAL GENERATIVE AI STUDY:

Business Rewards vs. Security Risks

Multi-sponsor report sponsored by:

protiviti®
Global Business Consulting

ISMG

TABLE OF CONTENTS

INTRODUCTION

Welcome to the **Second Annual Generative AI Study: Business Rewards vs. Security Risks**.

This survey of more than 360 business and cybersecurity professionals conducted in Q3 2024 comprises responses from two cohorts, business leaders – including CIOs, board members, executives or other business leaders – and CISOs or other cybersecurity professionals. Both groups represent a wide range of vertical sectors from around the world, with the largest group coming from North America.

In the survey, we look at the extent to which generative AI is deployed. Where it is deployed, we look at measuring perceived productivity gains, and where it is not currently used, we look at the anticipated benefits and intended deployment. This includes current and intended allocation of budget and projected growth as well as areas for investment going forward. We also look at the sometimes contrasting perspectives between business leaders and cybersecurity professionals when it comes to their current and intended use cases for generative AI.

We also compare prioritization of concerns, what those concerns are for each group, where they align and where they differ. We consider what mitigation strategies are being used or could be deployed to address these concerns.

The survey seeks to get a snapshot of respondents' understanding of current regulations.

More than just survey results, this report offers expert analysis of what organizations perceive to be the main security challenges and business opportunities associated with the deployment of generative AI. This report benchmarks what your competitors are doing so that you can use these results to help enhance your defenses and identify some of the productivity opportunities that generative AI presents.

While the first survey was characterized by hopes and fears, the second survey reflects the real-world successes and challenges of generative AI implementation.

TONY MORBIN

Executive News Editor, EU
Information Security Media Group
tmorbin@ismg.io

BY THE NUMBERS

Significant statistics that jump out from the Second Annual Generative AI Study: Business Rewards vs. Security Risks:

52%

of respondents have specific plans to purchase AI solutions over the next year.

AI use in production doubles from

15% to 31%

in a year.

Those not planning to use AI fell by more than three-quarters, from

27% to 6%

in a year.

There is a huge discrepancy in perception as

16%

of business leaders say AI is in production, but

34%

of security leaders say it is deployed.

Organizations with specific budgets for generative AI solutions double from

13% to 27%.

EXECUTIVE SUMMARY

Diverging Perspectives on Generative AI Deployment

The reported deployment of generative AI in production has risen from 15% to 31% in the second year after the public launch of ChatGPT - but that average masks a stark difference in perception about the extent of deployment between security leaders and business leaders.

Although 34% of security leaders in our survey report that generative AI is currently deployed in production in their organization, 16% of business leaders report such deployment. This statistic alone, on the fundamental issue of whether generative AI is even deployed in production, highlights the disparity in perception between security and business leaders when it comes to AI deployment. In many instances, this disconnect is echoed with different approaches and prioritization of use cases, security concerns, business opportunities and investments.

This initial mismatch suggests that either deployment is greater than business leaders realize, that definitions of what constitutes generative AI deployment differ between the two groups, and/or that business leaders' directives are either not being communicated or are ignored. Consequently, it is not a surprise that in addition to differences reflecting the remit of the two groups, their approaches vary in other aspects too.

The first survey in this series, conducted in Q3 2023, found business leaders were more enthusiastic than security leaders about implementing generative AI. They envisioned more use cases, were more optimistic about the potential gains and had fewer concerns about the risks of deployment.

In this 2024 survey, security leaders remain more cautious than business leaders, citing more concerns around the complexity of implementation, the difficulties of integration, the lack of skilled staff, the diversity of use cases, and the challenge of measuring, let alone achieving, ROI from generative AI deployment. While business leaders do share more success stories in terms of ROI and efficiency gains, they too are increasingly aware of the potential security and deployment issues - although both groups report lower levels of security issue concerns than a year ago.

Most notably, there has been a three-quarters fall in those saying they do not implement generative AI and have no plans to do so.

SURVEY RESULTS

01. To what extent does your company currently utilize generative AI?

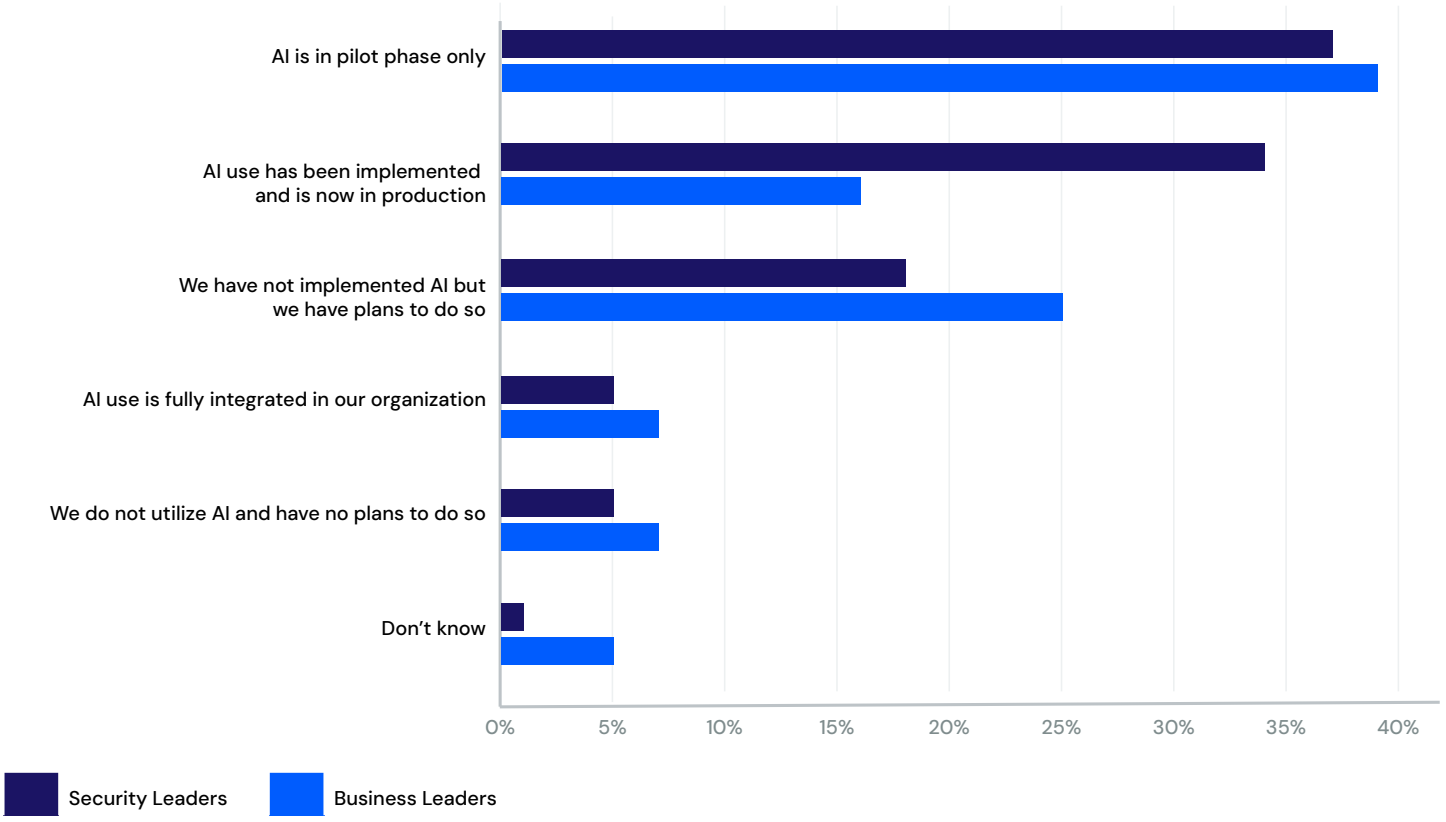

Figure 1: Generative AI Is Here - But Are You on the Same Page?

When it comes to generative AI, security and business leaders don't seem to agree on how far along their organizations are. While 34% of security leaders say AI is fully deployed in production, only 16% of business leaders think that's the case. This gap highlights a serious need for better communication and alignment. Are business leaders unaware of deployments, or do security leaders have a broader definition of what counts as "production"? Either way, it's a conversation that can't wait.

More than a year after the widespread availability of generative AI, marked by the launch of ChatGPT, the percentage of respondents saying AI use has been implemented and is now in production has risen from 15% to 25%. However, this year a further 6% say AI use is now fully integrated into the organization - something that was not reported a year ago.

When combining those who have generative AI fully integrated - 6% - with those who report generative AI use as now being in production - 25% - we get a figure of 31% of respondents actively using generative AI in production. This represents nearly a third of respondents - more than double the 15% reported a year ago.

Once again the largest group of respondents, 38%, say that AI is in the pilot phase only - an increase from 28% who reported AI in the pilot phase a year earlier.

Twenty-two percent of respondents say that they have not implemented AI but have plans to do so, which is at par with the 27% who said the same thing a year ago.

Meanwhile, 6% remain firmly in the skeptic/cautious camp, reporting, "We do not utilize AI and have no plans to do so," which is less than a quarter of the 27% who said the same a year ago.

Thus, we see a big shift toward AI acceptance and implementation and a reduction in AI skepticism/caution, but even this doubling in the transition to AI-driven production is still not as fast as many predicted. Part of the reason is explained later on where respondents describe implementation as more difficult and complex than anticipated, with huge variability on ROI dependent on use cases. While not broken out in this survey, it has been observed that the public sector, especially the health and finance sectors, is the most cautious about adopting AI and is most likely to have bans.

The biggest discrepancy between the two groups of respondents was in generative AI use in production: 34% of security leaders say that AI use has been implemented and is now in production, compared to 16% of business leaders saying the same. This is a reversal from a year ago when 17% of business leaders versus 13% of security leaders said that AI had been implemented and was in production. While both groups have shown an increase, the business leaders did so by a far smaller margin, and the perception gap between the two has widened but swung in favor of the security teams being more positive.

However, 5% of security leaders say generative AI use is fully integrated, compared to 7% of business leaders. Combining those who have AI fully integrated with those who report generative AI use now being in production, we find that generative AI is viewed as already implemented by 39% of security leaders but only by 23% of business leaders.

There was broad agreement regarding those in the pilot phase only, with 37% of security leaders and 39% of business leaders indicating they were in the pilot phase.

Some 18% of security leaders said that although generative AI is not implemented, there are plans to do so, compared to 25% of business leaders saying the same.

When it came to those not using generative AI and having no plans to do so, 5% of security leaders put themselves in this group, whereas 7% of business leaders said the same.

"Don't knows" were 1% among security leaders and 5% among business leaders.

Apart from the very top end of those having AI fully integrated, business leaders generally are less likely to believe their organization's adoption of AI is more progressed than those in security. Given the latter group is more likely to be operationally involved in implementation, it seems business leaders are surprisingly underestimating deployment – or are unaware of its extent.

When it comes to outright rejection/bans, security leaders are twice as likely as business leaders to believe this is the case. Since such decisions are likely made at the business leadership level, it seems the message is not being effectively communicated to security teams.

Best practice in security is to get an accurate assessment of the current situation before making decisions to improve, but when it comes to generative AI deployment, there is a difference in understanding of the current situation between security and business leaders over where that starting place might be. This suggests a more thorough monitoring and auditing of deployment should take place.

In this survey, no distinction is drawn between officially sanctioned deployment of AI and unsanctioned individual use via shadow AI, and it may be that this also plays a part in the differences in perception between business and security leaders.

02. Who in your organization is responsible for deploying generative AI productivity solutions (job title or role)?

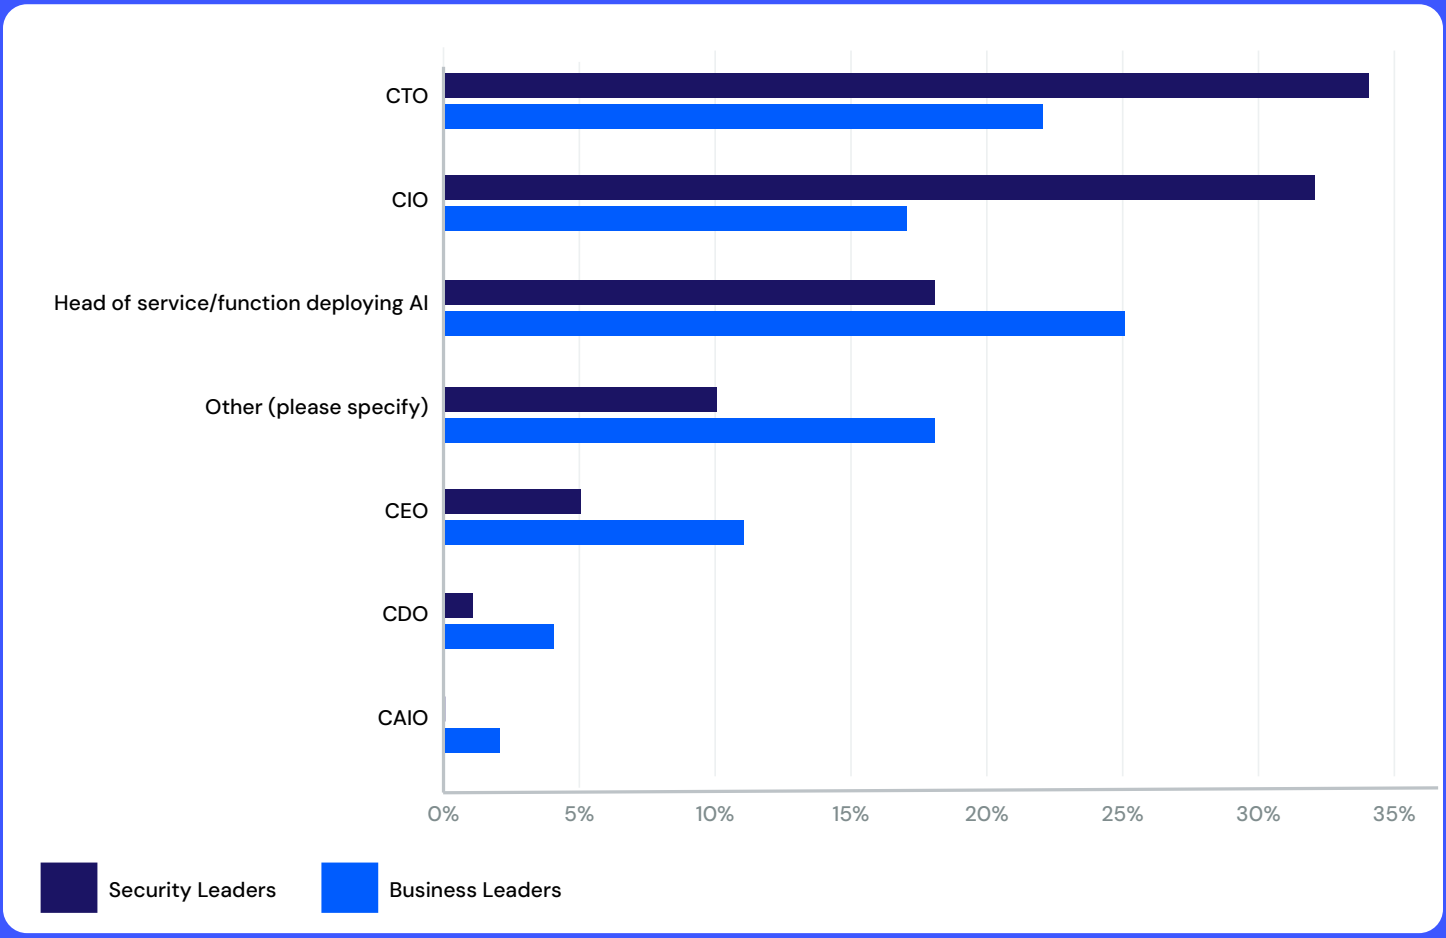

Figure 2: Who's Calling the Shots on Generative AI?

Who owns generative AI in your organization? Security leaders overwhelmingly point to CISOs and CTOs, while business leaders see responsibility spread across CIOs and other key roles. This disconnect could be a bottleneck for progress. If no one knows who's driving the bus, how can you steer AI strategies toward success? It's time to clarify roles and put the right leaders in the driver's seat.

There appears to be a technology orientation trend in AI responsibility, with the CTO being the most mentioned title at 27%, followed by the CIO at 25%. The third most popular title is head of service/ function deploying AI, at 21%.

The CEO continues to take responsibility for 8% of respondents, with 14% citing other titles.

A year ago, dozens of titles were being identified as being responsible for AI, suggesting a lack of consensus on the right person for the role. Some respondents stated that no one was responsible.

Among the named roles were CIO, CISO and CTO, with tech titles accounting for around two-thirds of responses, while head of service equivalents were just a handful.

This suggests that the direction of travel is toward more clearly defined roles for AI, from a tech-led activity to a line-of-business activity, even though it currently remains tech-dominated.

When comparing the responses from security leaders and business leaders, the latter were most likely to say that the head of service/function - at 25% - is responsible for deploying generative AI productivity solutions, followed by CTO at 22%, CIO at 17% and CEO at 11%.

In contrast, security leaders say the CIO at 32% or CTO at 34% is responsible for deployment, with just 18% identifying the head of service/function.

Given that business leaders decide who is responsible for what, it appears that this information is not effectively filtering down to security leaders. This impression of AI increasingly becoming a line of business-led activity may be correct in theory but overstated in practice. It suggests that in practice, the line of business is more likely to be actively involved in AI deployment in a hands-on manner than management perceptions or directives report.

03. Who in your organization is responsible for securing generative AI productivity solutions (job title or role)?

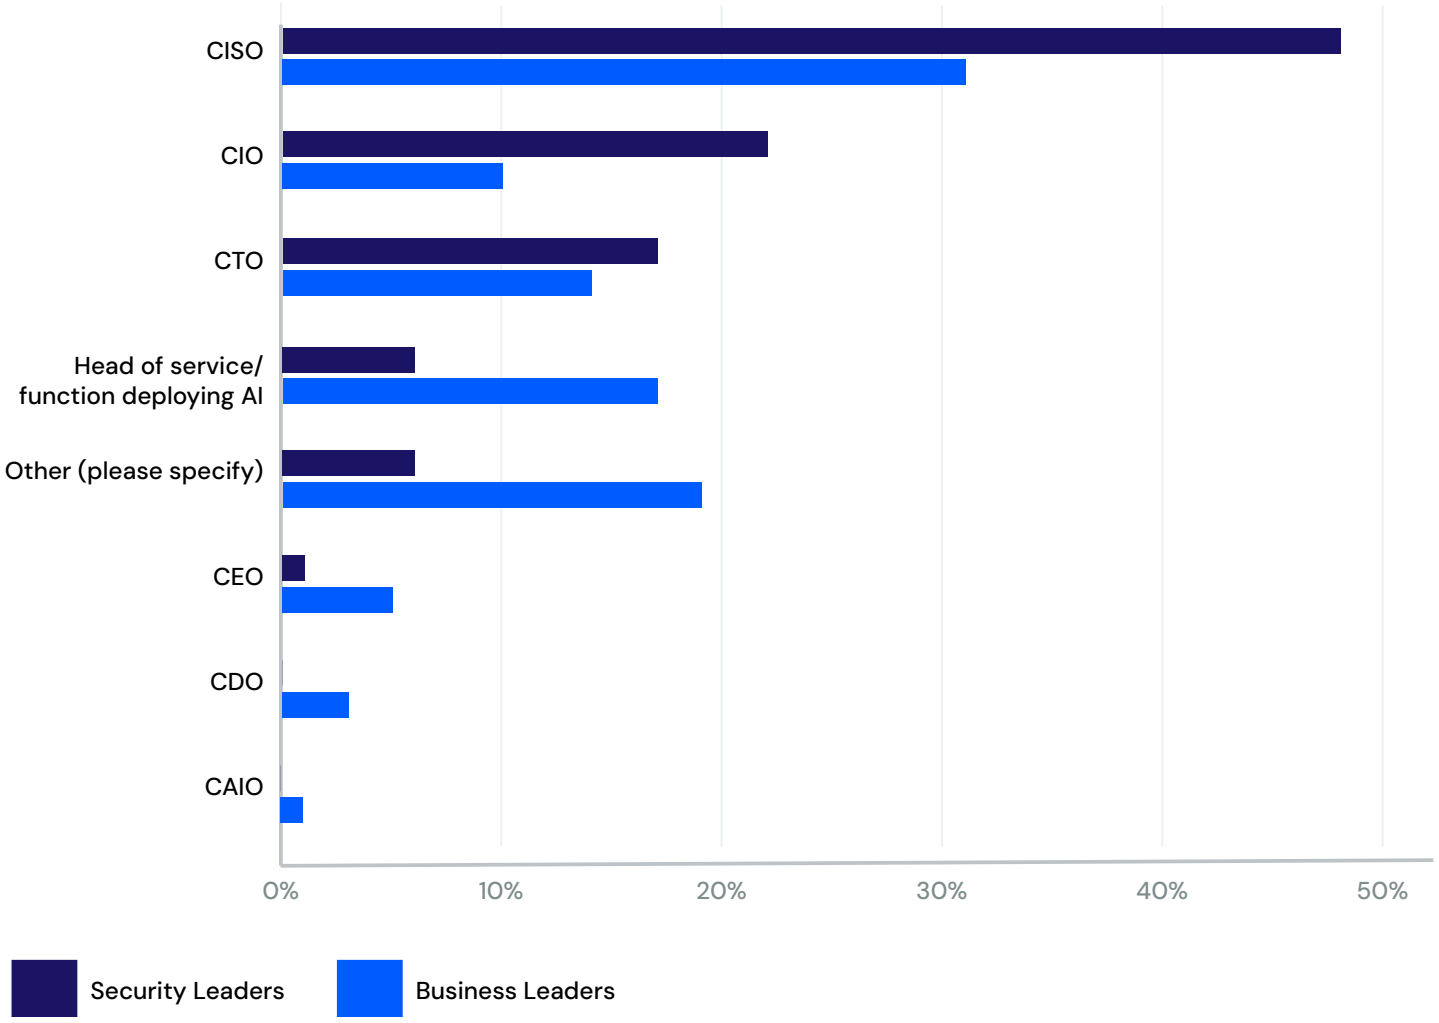

Figure 3: Generative AI Security - CISOs Take the Lead

According to security leaders, CISOs are taking charge of generative AI security, with nearly half - 48% - identifying themselves as the go-to person for AI security strategy. Business leaders, however, see security responsibility as more widely distributed.

The overwhelmingly popular title for securing generative AI productivity solutions is CISO at 39%, which is more than twice as popular as the second and third most-cited titles of CIO at 16% and CTO at 15%. The head of service/function deploying AI was at 11%.

CEOs garnered 3% of responses and a significant 13% cited other titles.

While this is similar to the results from a year ago, the role of the CISO in this regard has consolidated, and the number of "other" titles appears to have fallen (stats are not directly comparable for this question).

When analyzing the two groups, security leaders put CISO as the top answer at 48%. Although business leaders gave CISOs a lower percentage at 31%, it was still their top choice.

For security leaders, the CIO came in second at 22%, the CTO third at 17%, before dropping down to 6% saying head of service/function.

However, business leaders put the head of service/function in second place at 17%, followed by CTO at 14% then CIO at 10%.

While business leaders did put the CISO as the most likely to be in charge of security, they had a much higher rating for line of business as being responsible for securing generative AI.

04. If you use or plan to use generative AI, what are the primary strategic objectives you aim to achieve with generative AI within the next 2 to 3 years? [Respondents ranked issues as the most critical/important (7); very important (6), important (5), useful (4); and like to have (3, 2 and 1). The chart below shows only those responses scoring 7.]

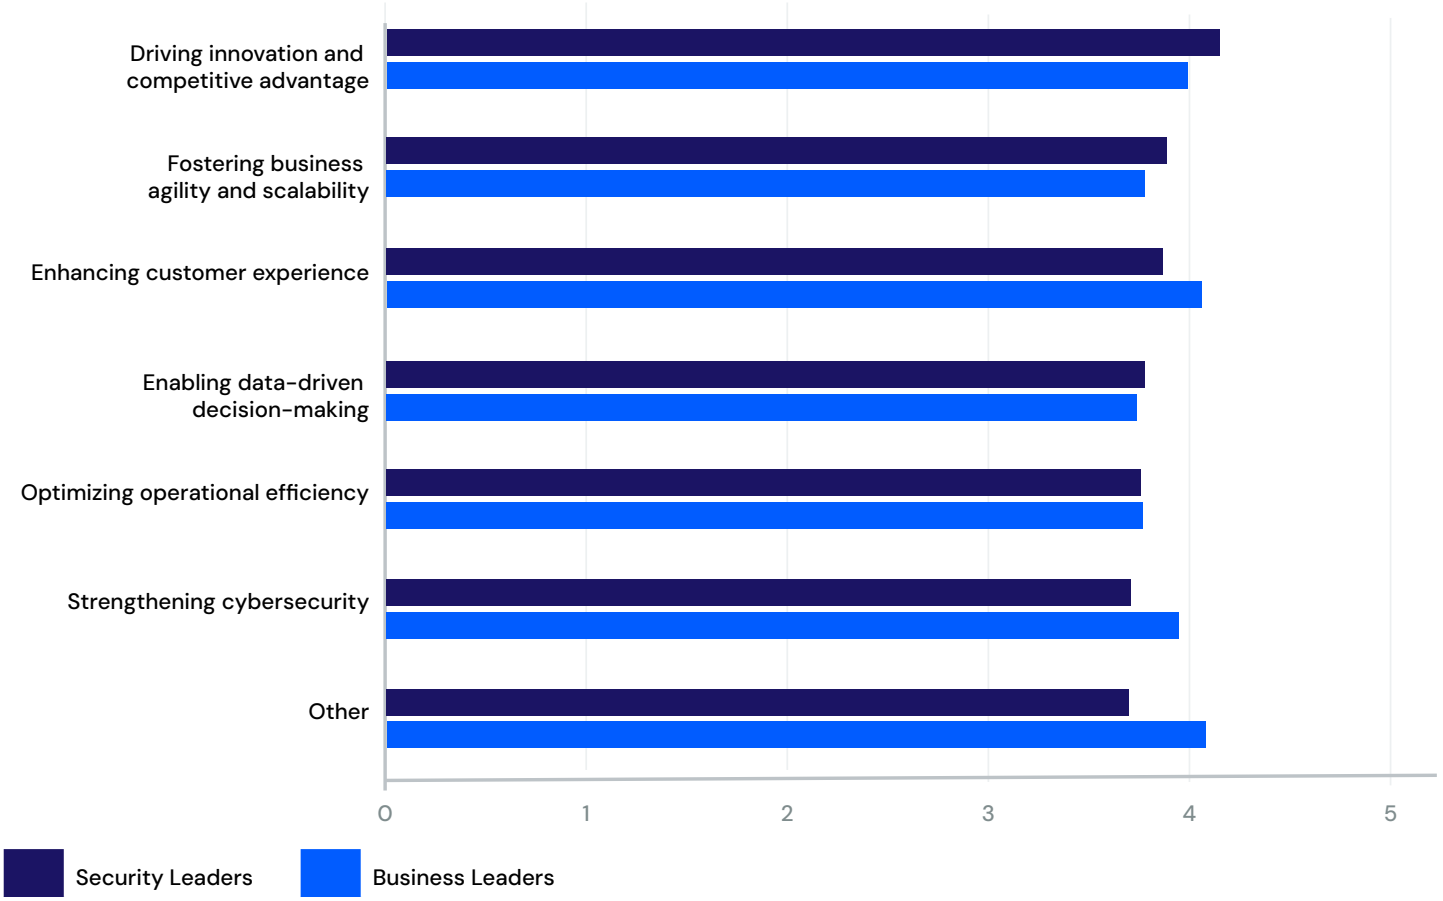

Figure 4: Why Are You Investing in Generative AI?

Generative AI offers big promises, but the reasons for investing vary depending on who you ask. Security leaders prioritize innovation and competitive advantage, while business leaders focus on improving customer experience. This difference in priorities shows how generative AI isn't a one-size-fits-all solution - it's about aligning strategies to get the most bang for your buck. What's your top priority?

There is a wide range of strategic objectives behind implementing AI, and when looking at issues rated the most critical/important (i.e., scoring 7), the "other" category of issues not specified by the survey is the most popular at 45%, ranking highest.

Of the issues suggested, the most important issue is "optimizing operational efficiency," with 15% of respondents giving it the highest ranking of 7, followed by "enhancing customer experience" at 14%.

"Driving innovation and competitive advantage" was tied with "strengthening cybersecurity," both getting the highest ranking score from 11% of respondents.

"Enabling data-driven decision-making" came in at 9% for the highest ranking score, while "fostering business agility and scalability" came in at 7% for the top score.

Looking at the top score only (i.e., those scoring 7 on the criticality scale), security leaders put "enhancing customer experience" top at 21%, followed by "optimizing operational efficiency" and "strengthening cybersecurity" at 15%, and "driving innovation and competitive advantage" at 9%.

For business leaders, there was less differentiation between their most critical issues, but the top place went to "optimizing operational efficiency" at 14%, followed by "driving innovation and competitive advantage" and "enhancing customer experience," both at 12% - marginally ahead of "strengthening cybersecurity" at 11%.

It is understandable that business leaders should prioritize "optimizing operational efficiency," and although even more security leaders rated it critical, it is nonetheless interesting that it did not lead their concerns, potentially suggesting misaligned priorities between business and security.

05. What are the main use cases you have/envision your organization implementing using generative AI or discriminative AI?

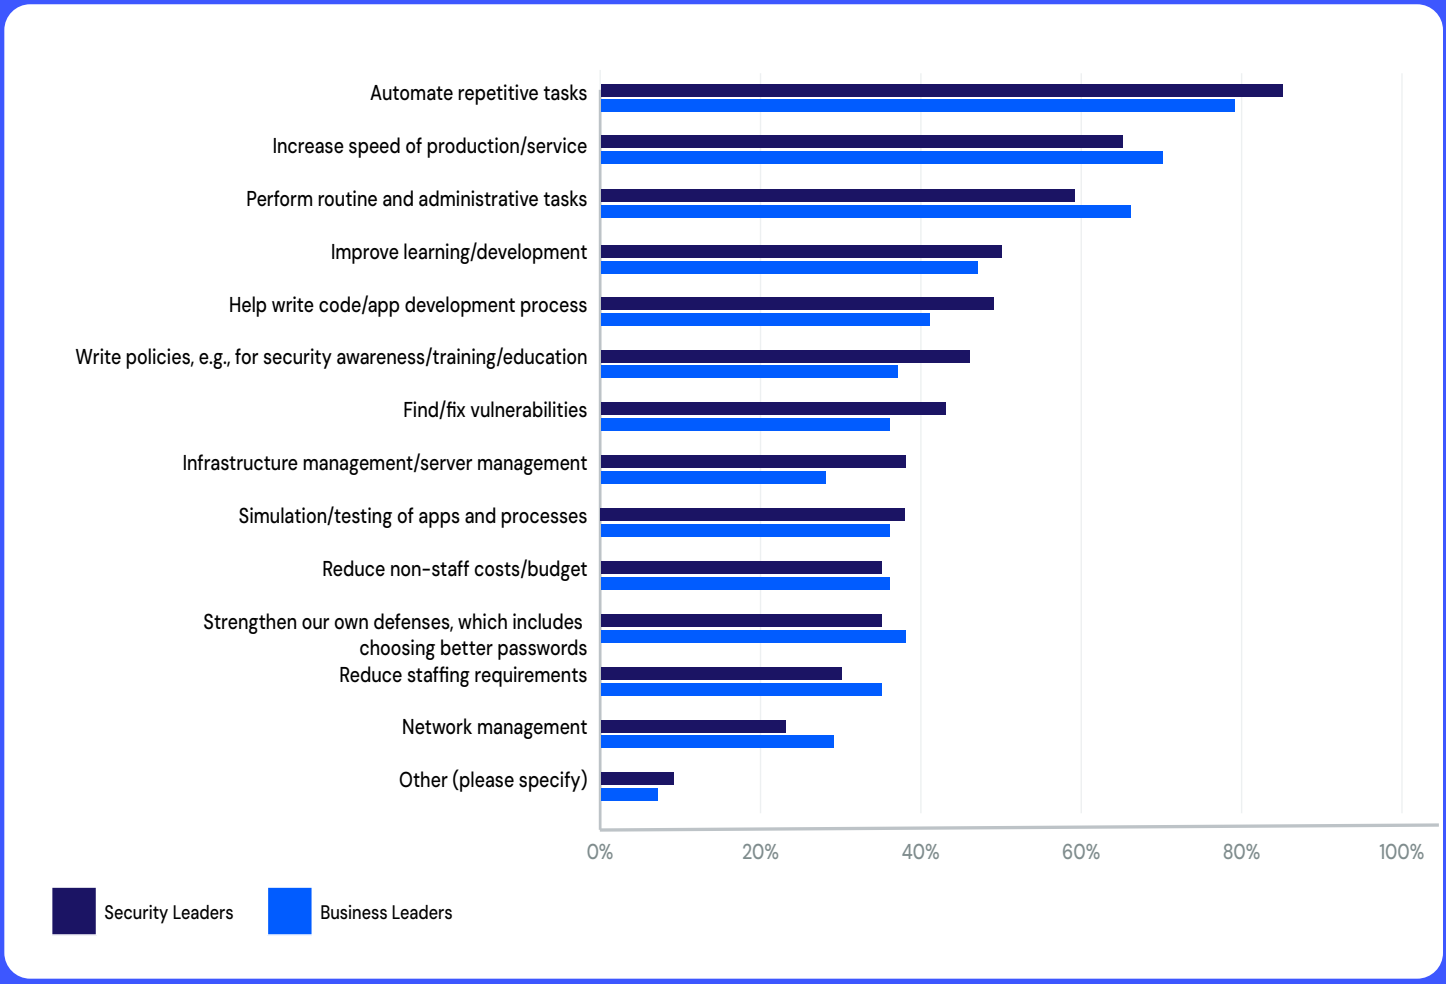

Figure 5: Generative AI Use Cases That Get Things Done

Automating repetitive tasks and speeding up production are the top reasons organizations are betting on generative AI. Security and business leaders mostly agree on these priorities, but there's a small divide when it comes to reducing staffing needs. Security leaders seem less focused on cutting headcount. The takeaway? Generative AI is all about efficiency – whether it's streamlining processes or boosting productivity.

The primary use case for generative AI or discriminative AI is to "automate repetitive tasks" at 82%, which is significantly up from the year-ago figure of 62%. This is followed by "increase speed of production/service" at 67%, slightly up from the year-ago figure of 59%. These two issues were also the leading priorities a year ago, though their importance has further increased.

In the third place is "perform routine and administrative tasks" at 62%.

Just below the halfway mark was "improve learning/development" at 49%, followed by "help write code/app development process" at 45%. "Write policies, for example, for security awareness/training/education" was cited by 42% and "find/fix vulnerabilities" by 39%. "Simulation/testing of apps and processes" followed closely at 37%.

Surprisingly, cost cutting was not a leading use case, but was still sought by a third of respondents: "Reduce non-staff costs/budget" at 36% and "reduce staffing requirements" at 32%, with the latter up from 24% a year ago.

While the previous year's survey is not directly comparable in this section, it still showed similar prioritization of concerns, particularly regarding automating repetitive tasks.

While there were differences in emphasis, for this question, business and security leaders were largely in agreement:

- Security leaders ranked "automate repetitive tasks" as the top use case at 85%, as did business leaders, though at a somewhat lower 79%.

- Both groups put "increase speed of production/service" in second place, cited by 65% of security leaders and 70% of business leaders.

- "Perform routine and administrative tasks" ranked third for both groups: security leaders at 59% and business leaders at 66%.

Both groups put "improve learning/development" in fourth place at 50% for security and 47% for business, averaging 49% for both groups. "Help write code/app development" scored 45%, but was a little higher among security respondents - 49% - compared to business respondents - 41%.

Other criteria were broadly aligned between the two groups.

06. In which environments do you use/envision your organization using generative AI or discriminative AI?

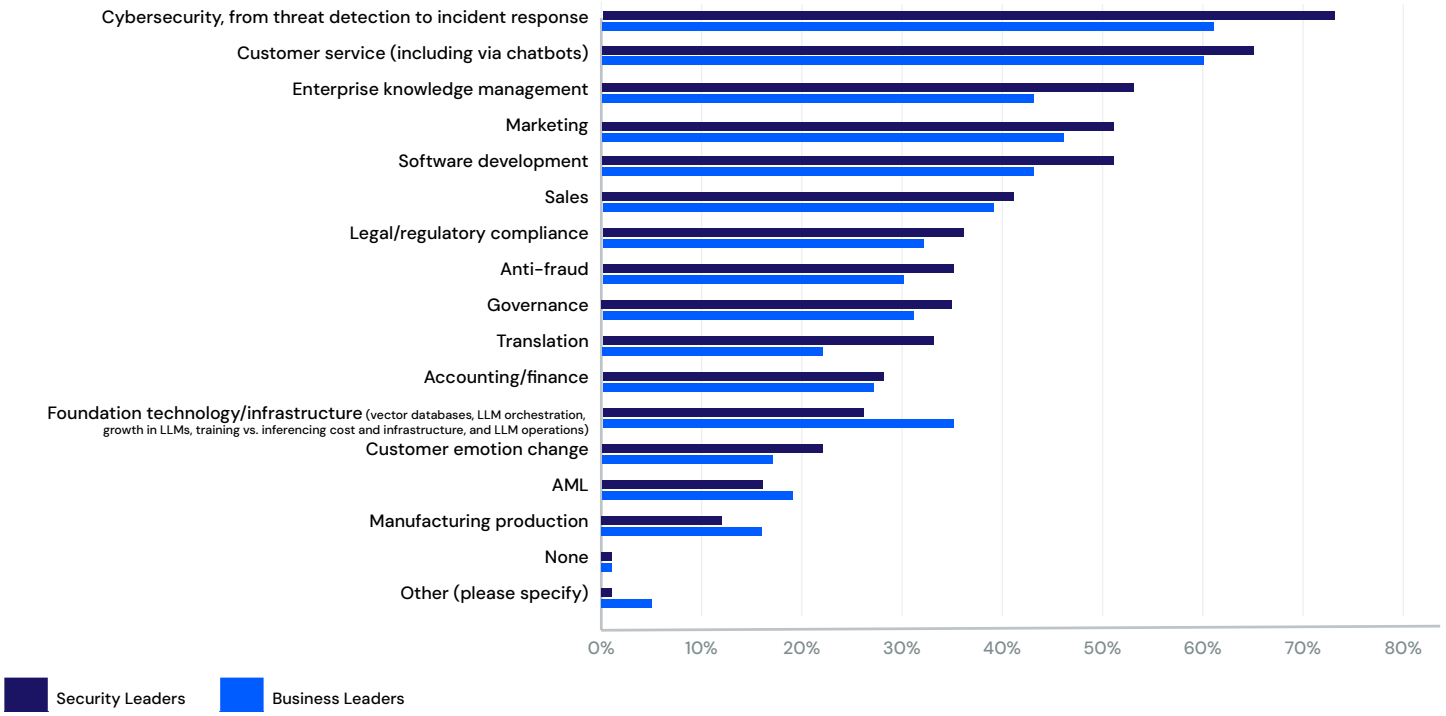

Figure 6: Generative AI – Beyond Back-Office Usage

While automation of back-office functions may be thought of as the low-hanging fruit for AI – and our figures show significant uptake for marketing and customer service bots – that is far from the whole story. In fact, use cases are led by cybersecurity, with enterprise knowledge management and software development also showing significant adoption.

"Cybersecurity, from threat detection to incident response" leads as the most anticipated environment for use at 67%, followed closely by "customer service (including via chatbots)" at 63%.

"Marketing" and "enterprise knowledge management," both at 48%, rank third, and "software development" just behind at 47%.

"Sales" leads the next level at 40%, followed by "legal and regulatory compliance" at 34% and governance at 33%.

There is broad alignment between security and business leaders as to which environments within the organization AI is used, although with some notable differences. Unsurprisingly, security leaders cite "cybersecurity, from threat detection to incident response" to a larger degree, at 73%, compared to business leaders of whom 61% rank it first. "Customer service (including via chatbots)" comes second, with 65% of security leaders and 60% of business leaders citing this use case.

Several other areas show significant divergence between the two groups, with "enterprise knowledge management" being cited by 53% of security leaders versus 43% of business leaders and "software development" by 51% of security leaders versus 43% of business leaders. "Foundation technology/ infrastructure" reveals an inverse trend, at 26% for security leaders and 35% for business leaders.

07. What challenges have you faced in integrating generative AI with existing IT infrastructure?

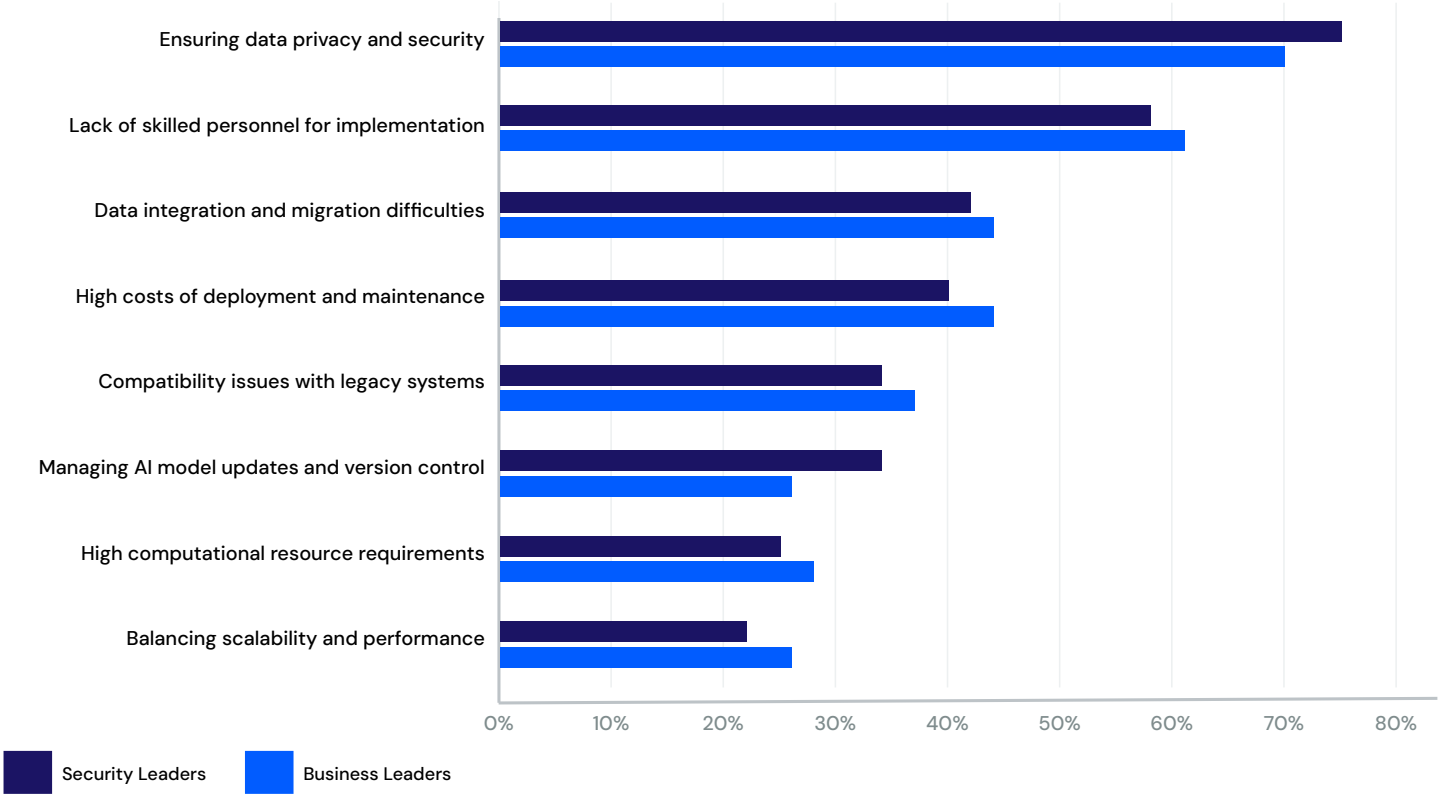

Figure 7: Challenges That Keep Generative AI Grounded

Generative AI comes with its own set of roadblocks, from implementation hurdles to a lack of skilled talent. Both security and business leaders are concerned about data privacy and security. With aligned priorities around these concerns, the teams must work together to tackle these challenges head-on.

"Ensuring data privacy and security" emerged as the primary challenge for integrating generative AI with existing IT infrastructure, with 73% of respondents highlighting this concern. This represents a slight decrease from the year-ago figure of 80% for "leakage of sensitive data by staff."

"Lack of skilled personnel for implementation" ranked second at 60%, underscoring a problem in this sector that was not mentioned in the previous survey.

"Data integration and migration difficulties" ranks third on the list of challenges at 43%, not mentioned previously. "High costs of deployment and maintenance" followed closely at 42%, with "compatibility issues with legacy systems" at 36%.

Although "ensuring data privacy and security" was ranked first by both security leaders and business leaders, security leaders rated it higher at 75% compared to business leaders at 70%. "Lack of skilled personnel for implementation" came second, with 58% for security leaders and 61% for business leaders. "Data integration and migration difficulties" ranked third for both security leaders - 42% - and business leaders - 44%.

Additional challenges included "compatibility issues with legacy systems," which scored 34% among security leaders and 37% among business leaders; "managing AI model updates and version control" received 34% from security leaders and 26% from business leaders; "high computational resource requirements" and "balancing scalability and performance" received 25% and 22% from security leaders and 28% and 26% from business leaders, respectively.

08. If you currently employ AI systems, what productivity gains do you estimate to achieve compared to the systems they replace?

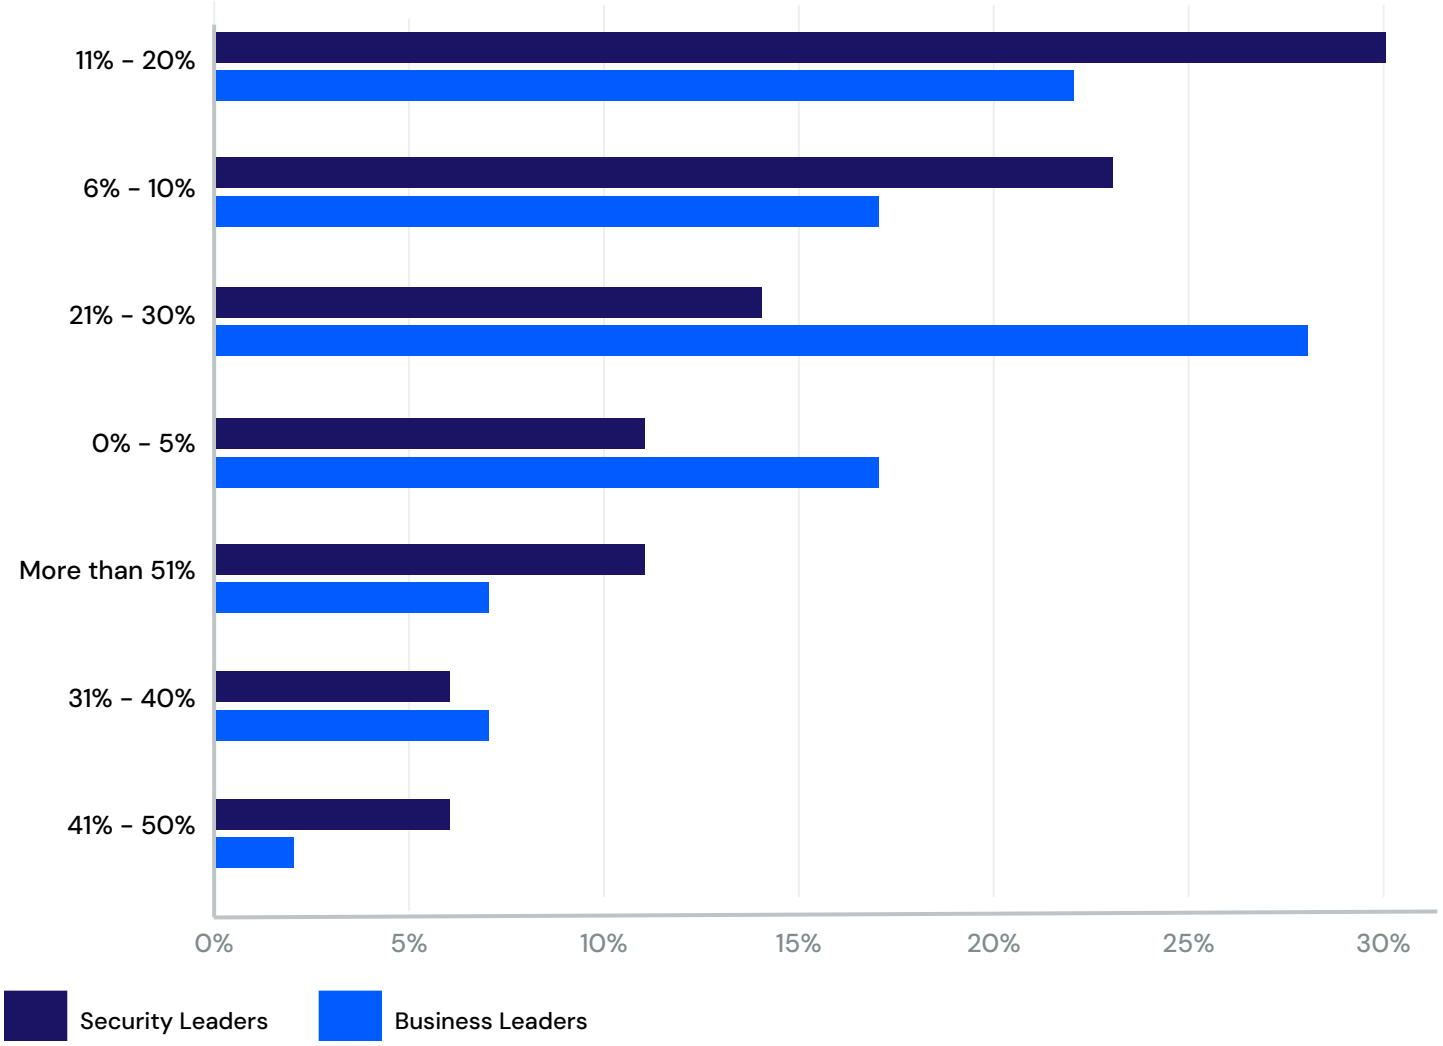

Figure 8: AI and ROI - Are We There Yet?

Everyone wants to know if generative AI will pay off. While business leaders report early success stories, security leaders remain skeptical, citing challenges in measuring impact. This figure sheds light on how organizations are approaching ROI and what it means for their AI strategies.

Forecasts of 11% to 20% gains were cited by 26% of respondents, slightly ahead of those respondents forecasting 21% to 30% gains in productivity - 21% of respondents - with a further 20% forecasting gains of 6% to 10%. The majority of respondents - some 67% - expect the gains to be between 6% and 30%.

At the most pessimistic end of the scale, 14% of respondents forecast gains of less than 5%, while at the more optimistic end of the scale, all those forecasting more than 30% gains totaled 20%, thus one

in five. This included a highly optimistic 9% of respondents who forecast gains of more than 50%.

Overall, it shows significant optimism around the potential productivity gains to be achieved. When comparing with the previous year's figures, it is difficult to judge whether the early adopters in last year's report were achieving gains based on targeting the low-hanging fruit, and are now seeing less substantial gains, or that more respondents are in a position to judge based on more robust use cases and maturing generative AI models.

Security leaders and business leaders hold diverse opinions regarding their optimism about productivity gains expected from AI.

At the lowest level, it initially appears as if business leaders are more pessimistic with just 11% of security leaders predicting the lowest score of 0% to 5%, compared to 17% of business leaders.

However, the 6% to 10% category is selected by 23% of security leaders and 17% of business leaders, the 11% to 20% category is the most popular answer for security leaders at 30%, compared to 22% for business leaders. But 21% to 30% gains, selected by 14% of security leaders, is the top answer for business leaders at 28%.

Although these figures show a mixed pattern, it does appear that business leaders are both more likely to be pessimistic or optimistic while security leaders are largely in the moderately optimistic range of between 6% to 20% accounting for over half of their responses - 53% - whereas the largest grouping for business responses was slightly more optimistic, with 11% to 30% gains forecast by 50% of business leaders.

09. How do you measure the return on investment (ROI) for generative AI projects?

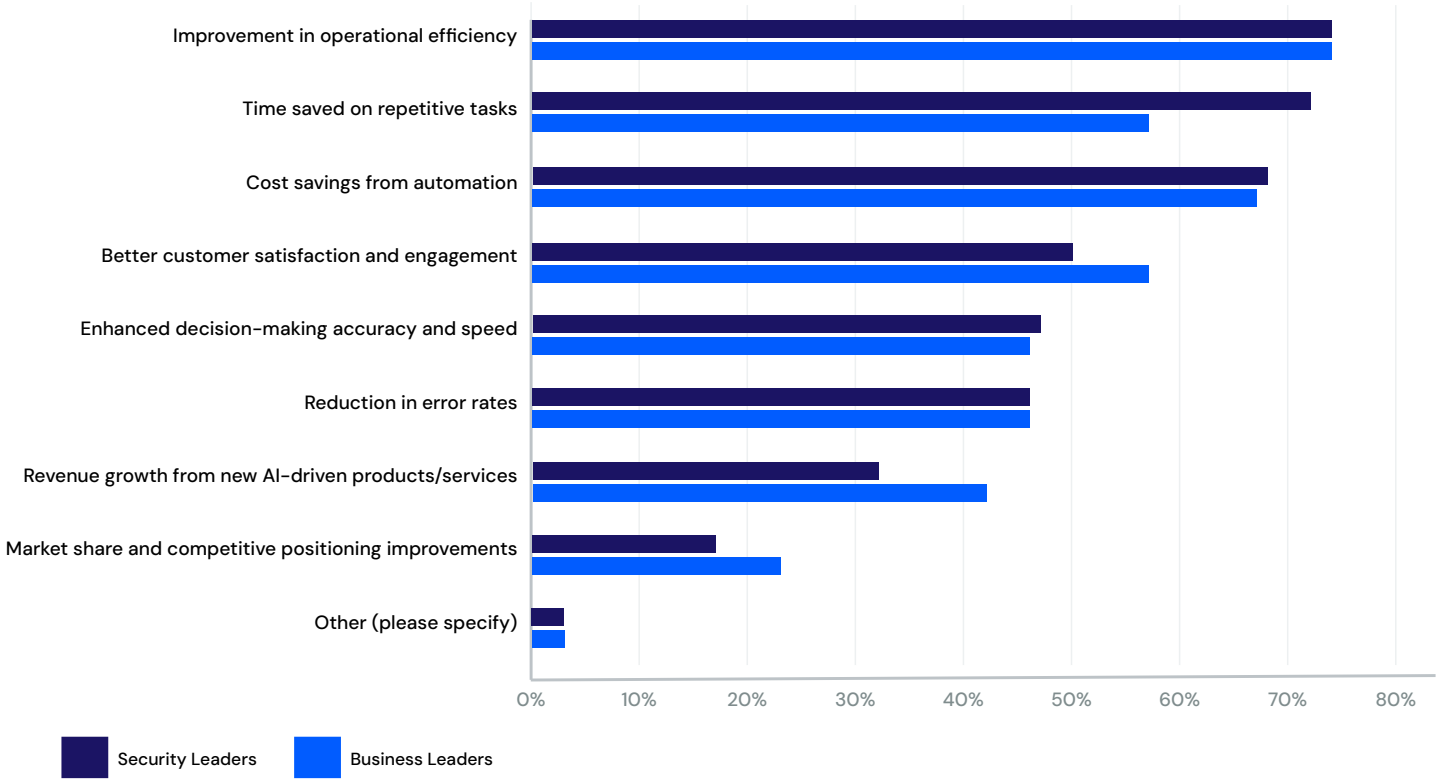

Figure 9: Generative AI in Action - What's Working?

From automating customer service to optimizing workflows, organizations are finding creative ways to use generative AI. However, not all use cases are created equal. This figure highlights the most successful applications and where organizations are struggling to find value.

"Improvement in operational efficiency" emerged as the primary ROI measure, cited by 74% of respondents, closely followed by "cost savings from automation" at 67%. "Time saved on repetitive tasks" ranked third at 65%, while "better customer satisfaction and engagement" claimed the fourth position at 53%.

Improvements in other aspects of service provision also received significant support, with "enhanced decision-making accuracy and speed" scoring 47% and "reduction in error rates" at 46%. These are followed by "revenue growth from new AI-driven products/services" at 37% while "market share and competitive positioning improvements" was cited by 20%.

When defining the ROI for generative AI, both security leaders and business leaders aligned closely on "improvement in operation efficiency" and ranked it their top priority at 74% and 73%, respectively.

Security leaders put "time saved on repetitive tasks" in second place at 72%, whereas this was scored at 57% by business leaders. Business leaders gave second place to "cost savings from automation" at 67%, which received a higher score from security leaders at 68% but a lower ranking.

Security leaders scored "better customer satisfaction and engagement" and "enhanced decision-making accuracy and speed" at 50% and 47%, respectively, while business leaders scored 57% and 46%, respectively.

However, "revenue growth from new AI-driven products/services" was cited by just 32% of security leaders compared to 42% of business leaders. Other scoring included "reduction in error rates" - scored 46% by both security and business leaders - while "market share and competitive positioning improvements" scored 17% by security leaders and 23% by business leaders.

10. Do you have a specific budget for generative AI solutions?

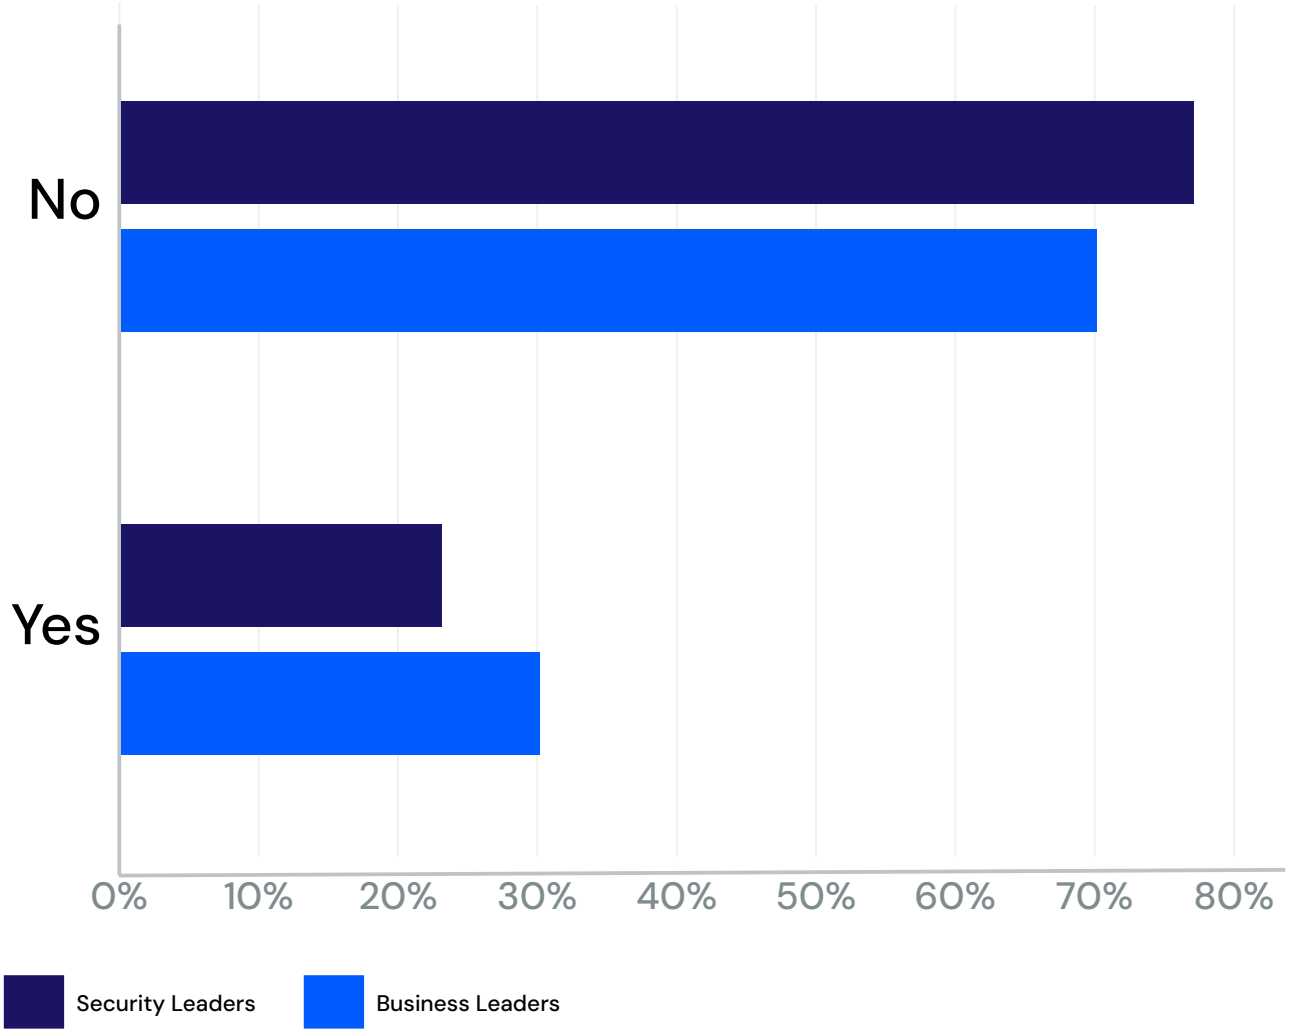

Figure 10: Budgets for AI – Who's Spending and Why

Generative AI budgets are growing, but who's managing the purse strings? This figure breaks down how security and business leaders are allocating resources to generative AI initiatives and what it means for the future of AI adoption.

In this year's survey, 27% of respondents say they have a specific budget for generative AI solutions. This is more than double the 13% saying the same a year ago, which is perhaps one of the more significant findings of the survey, confirming the rising importance of AI over the year.

Among the two groups, 23% of security leaders report a dedicated budget, compared to 30% of business leaders.

11. If no, do you expect to have one within 12 months?

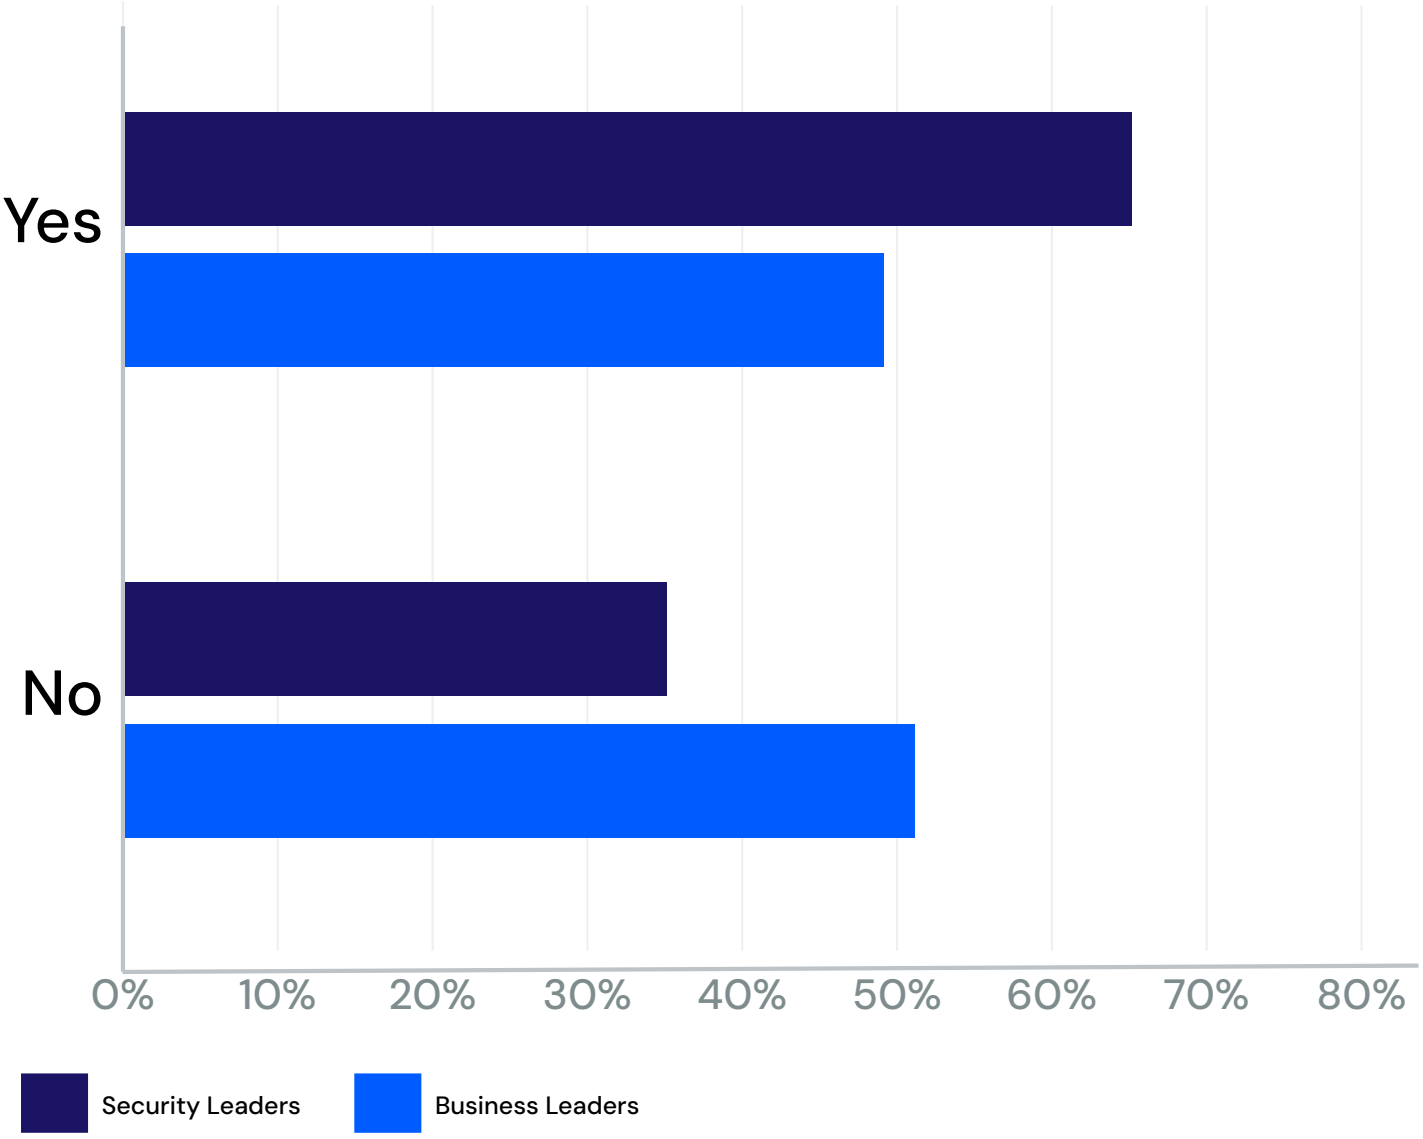

Figure 11: Generative AI and Security - A New Dedicated Budget Category

If all those expecting AI budgets in the table below - 42% of the total - are added to those currently with dedicated AI budgets in the previous table - 27% of the total - and if those expectations are delivered, 69% will have dedicated AI budgets in a year's time.

Among 73% of respondents who do not have a dedicated budget currently, 57% expect to have one within a year. A year ago, 54% expected this to be the case - but as can be seen from the earlier question, this did not always materialize.

So while 54% a year ago expected a dedicated budget adoption, the reality is that those with a dedicated budget actually increased from 13% to 27%; if the same doubling were to occur, the current 27% would rise to around 54%.

Some 65% of security leaders who did not have a dedicated AI budget expected to have one within 12 months, compared to 49% of business leaders.

12. If yes, what percent increase do you expect in 12 months' time?

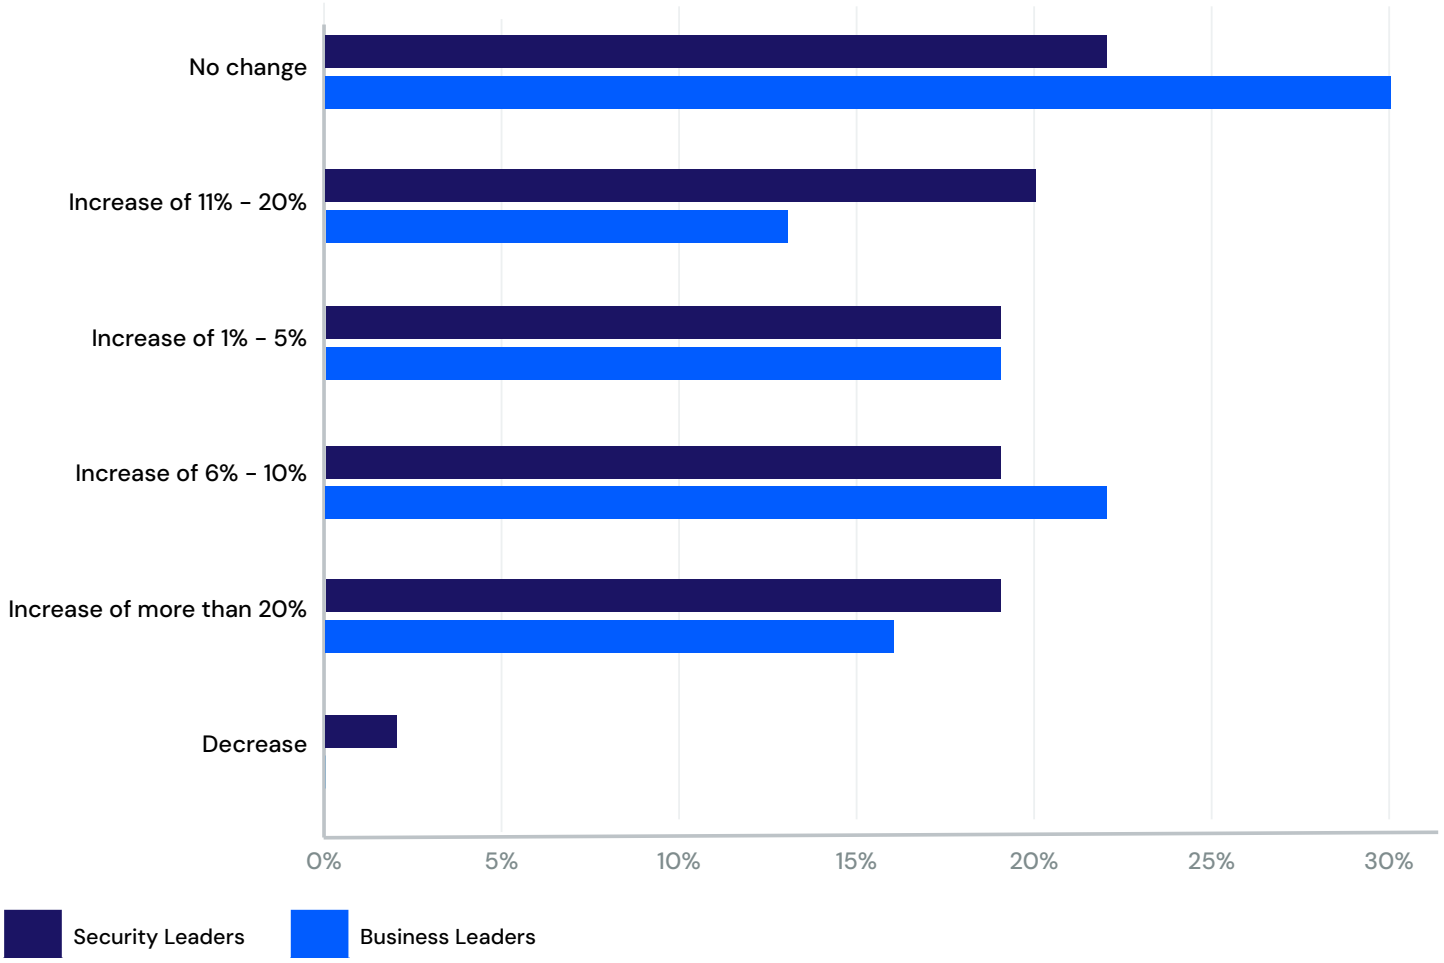

Figure 12: 70% of Respondents Expect Their AI Budgets to Increase

While exact forecasts vary, the overall picture is clear, which is that generative AI budgets are set to rise massively according to most respondents.

Of those who do have a dedicated generative AI budget, 26% expect to see no change in the coming year, the same proportion as a year ago, while less than 1% forecast a reduction. Thus, more than 70% are expecting a budget increase.

The largest group of these – at 20% – are expecting an increase of 6% to 10%; 19% report expected increases of under 5%; 17% forecast growth in their budget of 11% to 20%; and a further 18% forecast increases greater than 20% of the budget.

This represents a shift from last year when expectations were weighted more toward the lower end, indicating rising confidence in AI investments.

While 22% of security leaders expect no change, 30% of business leaders say the same. Although an increase of 6% to 10% was the most popular choice for business leaders, security leaders put that at 19%, with a more even spread across all the options.

13. Do you have specific plans to purchase AI solutions over the next 12 months for any of the use case options in Question 5 above?

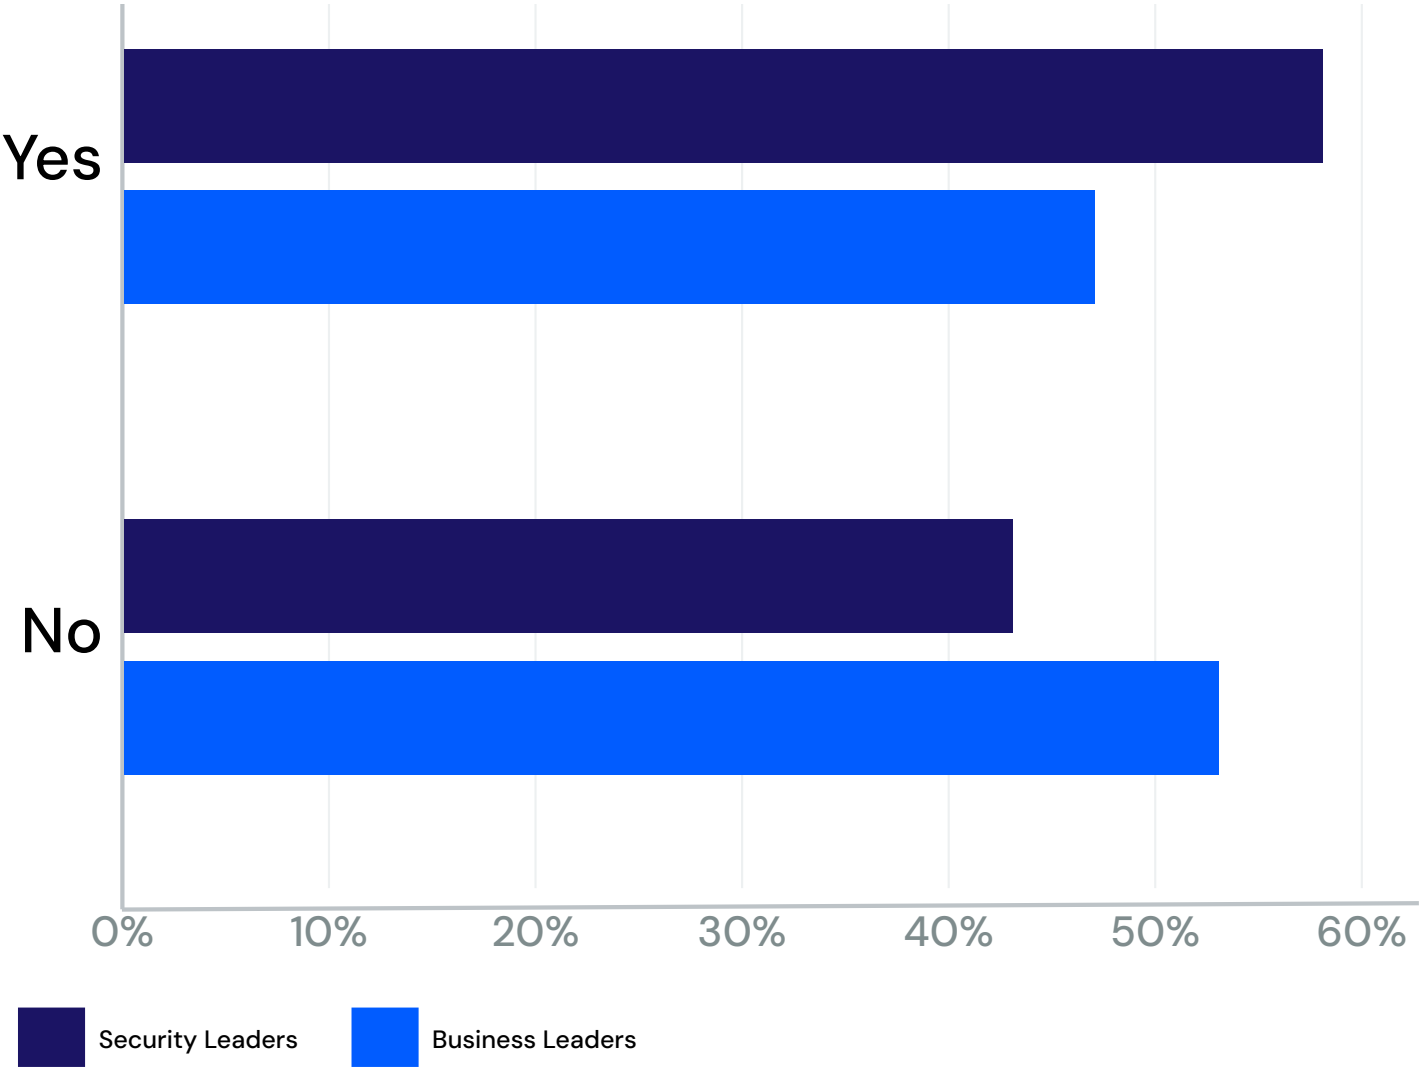

Figure 13: The Future of AI - What's Next?

As generative AI continues to evolve, organizations are looking ahead to specific deployments. This figure suggests concrete plans are now in place as to where organizations will implement AI solutions.

Just over half of respondents - 52% - say that they do have specific plans to purchase AI solutions over the next year. This compares to 31% saying the same a year ago. This increase, from 31% to 52%, represents 68% growth in the numbers with specific plans, likely reflecting both growing acceptance of AI technology and wider availability of more diverse models for more varied use cases.

More security leaders - 58% - than business leaders - 48% - say they had specific plans.

14. What are your main concerns when it comes to implementing generative AI?

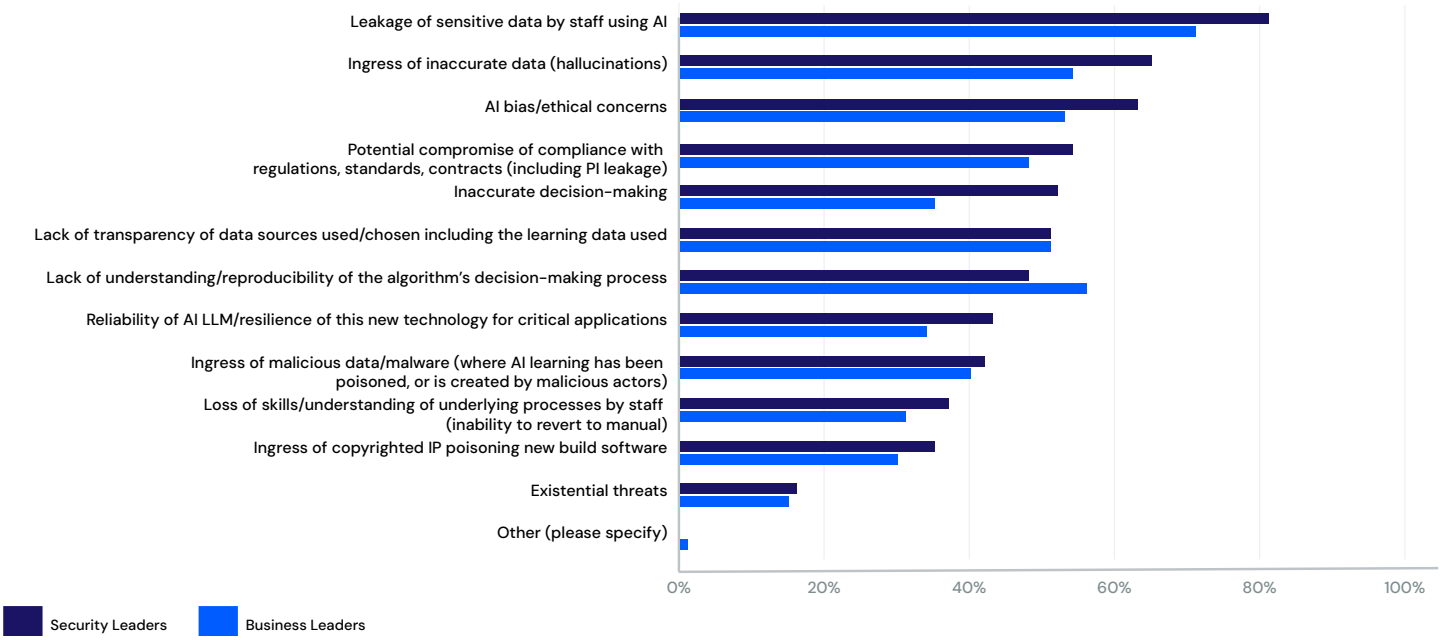

Figure 14: Generative AI – Risks Remain But Concerns Reduce

The potential risks associated with generative AI adoption are increasingly understood, but the level of concern surrounding these risks has significantly reduced as organizations experience actual implementation.

When ranking concerns about implementing generative AI, "leakage of sensitive data by staff using AI" emerged as the leading concern, cited as the top concern by 76% of respondents. This was also the top concern a year ago, with 81% citing it at the time, suggesting the concern, while still very important, may have slightly diminished.

The next highest ranked concerns are "ingress of inaccurate data (hallucinations)," mentioned by 59% of respondents, and "AI bias/ethical concerns" mentioned by 58% of respondents. Compared to a year ago, when these concerns were mentioned by 69% and 59% of respondents, respectively., There has been a decline in concerns about hallucinations, though they still remain significant, while ethical concerns remain remarkably consistent and still high.

The other issues mentioned by more than half of this cohort are "lack of understanding/ reproducibility of the algorithm's decision-making process" at 52% and "lack of transparency of data sources used/chosen including the learning data used" and "potential compromise of compliance with regulations, standards and contracts (including PI leakage)" both at 51%.

However, among other concerns getting significant responses are "inaccurate decision-making" - which scored 43% - "ingress of malicious data/malware (where AI learning has been poisoned or is created by malicious actors)" was cited by 41%, and "reliability of AI LLM/resilience of this new technology for critical applications" was cited by 38%.

"Loss of skills/understanding of underlying processes by staff (inability to revert to manual) was cited by 34%, "ingress of copyrighted IP poisoning new build software" by 32%, and "existential threats" by 15%.

These rankings broadly align with the previous year's findings, showing consistent patterns of concern.

"Leakage of sensitive data by staff using AI" was the top concern across both groups, but cited by more security leaders - 81% - than business leaders - 71%.

"Ingress of inaccurate data (hallucinations)" came second place for security at 65% but was third on the business leader ranking at 54%. Instead, security leaders' third placed "AI bias/ethical concerns" with 63% was fourth for business leaders at 53%.

Other results were a mix of agreement and disagreement, divided as follows: "potential compromise of compliance with regulations, standards, contracts (including PI leakage)" saw a divide with security at 54% and business at 48%. "Lack of understanding/reproducibility of the algorithm's decision-making process" was at 48% for security and 56% for business. However, "inaccurate decision-making" was one of the biggest divides, with security on 52% and business on 35%. "Lack of transparency of data sources used/chosen including the learning data used" was closely aligned, with both security and business on 51%.

"Reliability of AI LLM/resilience of this new technology for critical applications" again saw a wide divide with security at 43% and business at 34%. "Ingress of malicious data/malware (where AI learning has been poisoned, or is created by malicious actors)" saw a closer alignment, with security at 42% and business at 41%. There was more agreement also when it came to "loss of skills/understanding of underlying processes by staff (inability to revert to manual)" with security at 37% and business at 31%.

"Ingress of copyrighted IP poisoning new build software" was also close with security at 35% and business at 30%. While initially the figures for those concerned about "existential threats" may seem relatively low - with security at 16% and business at 15% - they highlight an important finding: about one in five knowledgeable industry professionals sees this new fundamental aspect of doing business as potentially creating a threat to humanity.

15. What tools, processes or approaches do you currently use to mitigate the concerns around use of AI by your own organization?

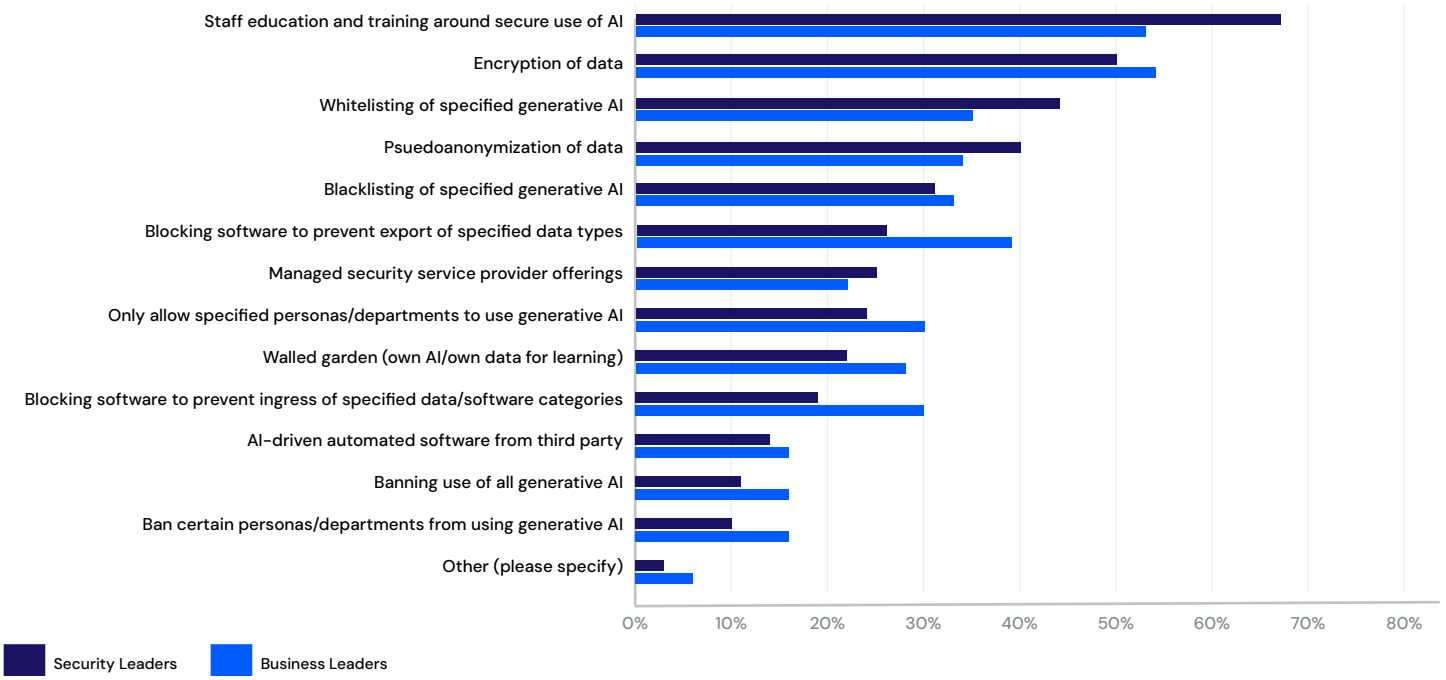

Figure 15: Risk Reduction – Mitigation Combines Education and Technology

Increasingly, organizations are able to identify mitigation strategies to reduce potential risks from generative AI implementation. These now combine staff education, technological solutions and processes deployed – though security and business leaders don't always agree on priorities.

Approaches to tackle these concerns are many and varied, but are led by "staff education and training around secure use of AI" at 60%, followed by "encryption of data" at 52%.

"Whitelisting of specified generative AI" is used by 40%; "psuedoanonymization of data" by 37%; "blocking software to prevent export of specified data types" by 33%; "blacklisting of specified generative AI" by 32%; and "only allow specified personas/departments to use generative AI" at 27%. "Walled garden (own AI/own data for learning)" and "blocking software to prevent ingress of specified data/software categories" are employed by 25% and 23% employ "managed security service provider offerings."

There are still 15% of respondents saying they use "AI-driven automated software from third party." Fourteen percent of respondents ban all use of all generative AI, while 27% only allow certain personas/departments to use generative AI, and 13% ban certain personas/departments from using AI. Security professionals were significantly more likely to rate "staff education and training around secure

use of AI" as the primary tool to mitigate their concerns around AI use, at 67%, compared to business leaders at 53%. However, business leaders ranked "encryption of data" as their first priority at 54%, while security leaders ranked it second at 50%.

Security professionals ranked "whitelisting of specified generative AI" as third at 44%, compared to 35% by business leaders.

Business leaders chose "blocking software to prevent export of specified data types" as the third most important approach - 39% - which had only been cited by 26% of security leaders.

"Psuedoanonymization of data" was cited by 40% of security leaders, compared to 34% of business leaders. This is followed by "blacklisting of specified generative AI, with security leaders scoring it 31% versus business leaders at 33%. "Managed security service provider offerings" scored 25% for security leaders and 22% for business leaders.

"Only allow specified personas/departments to use generative AI" was cited by 24% of security leaders and 30% of business leaders. "Ban certain personas/departments from using generative AI" came in at 11% for security and 16% for business, and "banning use of all generative AI" was 11% for security and 16% for business - thus the approach of using bans was slightly more favored by business than security.

Other approaches include "walled garden (own AI/own data for learning)" with security at 22% versus business at 28%; "blocking software to prevent ingress of specified data/software categories with security at 19% versus business at 30%; "AI-driven automated software from third party" was scored 14% by security leaders and 16% by business leaders.

It was interesting to see that the technical specialists - the security leaders - appear more likely to favor staff education, whereas business leaders favored technical solutions such as blocking software. This could be because the security leaders are more aware of the limitations of technology, or equally, it could be because the business leaders have taken a data-driven ROI approach.

16. Do you know and understand what regulatory restrictions or guidance applies to your use of generative AI in your geography/ industry vertical?

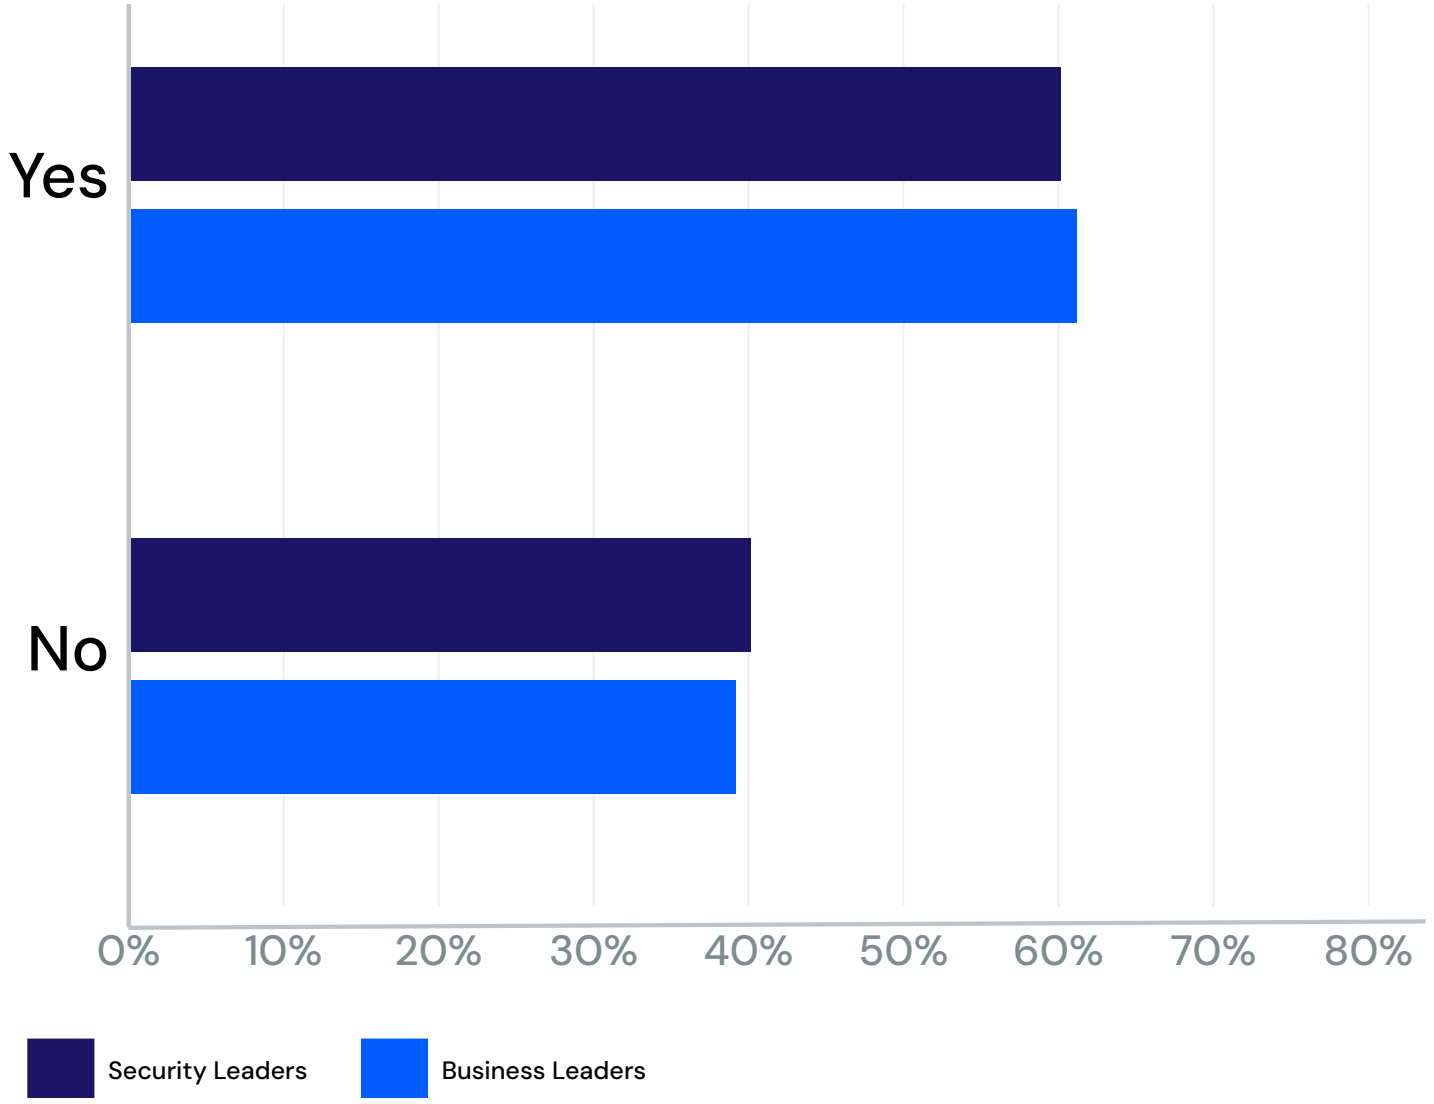

Figure 16: Regulations Increasingly Understood

While there remains a long way to go and the rapid pace of change in AI regulations can be hard to keep up with, our figures show significant growth in understanding of relevant regulations year on year.

Currently, 61% of respondents say they do know and understand what regulatory restrictions or guidance applies to their use of generative AI in their geography or industry vertical, up from 45% a year ago. This indicates a perceived increase in awareness and is likely, to some extent, to reflect an actual increase in understanding.

There was no significant difference between the two groups, with 60% of security leaders and 61% of business leaders reporting that they know and understand the relevant regulations in their sector.

While the rapid development of regulations in the sector explains the relatively low level of understanding, it is nonetheless a concern that some 40% of those responsible for implementing generative AI are not able to say that they understand the regulations with which they must comply.

17. How are you seeing cyber adversaries employ AI in their attacks against your organization?

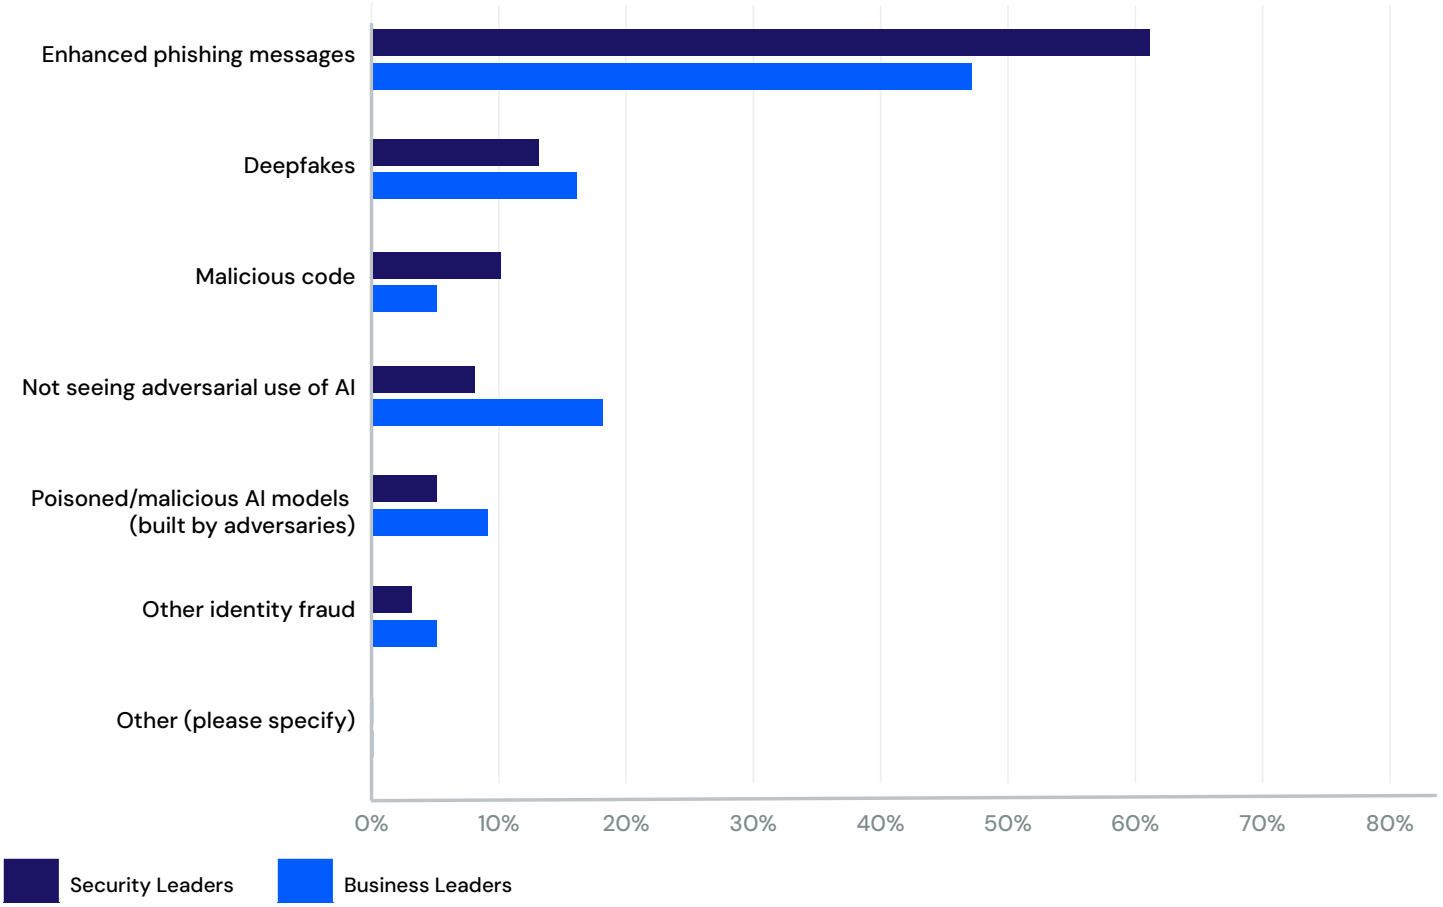

Figure 17: Adversarial Use of Generative AI

Generative AI is being used by adversaries, but it is not always seen. Nonetheless, most respondents were able to identify where attackers are now utilizing this new technology to enhance their own activities.

The primary use of AI by attackers currently being seen is for "enhanced phishing messages," reported by 54% of respondents. Other uses include "deepfakes" at 15%, "malicious code" at 8%; "poisoned/malicious AI models (built by adversaries)" at 7%; and "other identity fraud" at 4%.

However, 13% of respondents said that they are not seeing the adversarial use of AI.

When data from the two groups are broken out, the security leaders are more likely to report seeing AI used by adversaries for "enhanced phishing messages" at 61% compared to 47% of business leaders.

There is a divergence in the reporting of malicious code, with 10% of security leaders citing versus 5% of business leaders, and "deepfakes," with 13% of security versus 16% of business leaders reporting it.

However, security leaders were less likely to report "other identity fraud," with 3% observing it compared to 5% of business leaders, and "poisoned/malicious AI models (built by adversaries)," with 5% reporting this threat compared to 9% of business leaders.

Security leaders were also less likely to say they were "not seeing adversarial use of AI" at 8% compared to business leaders at 18%.

18. How are you addressing the skills gap related to generative AI within your organization?

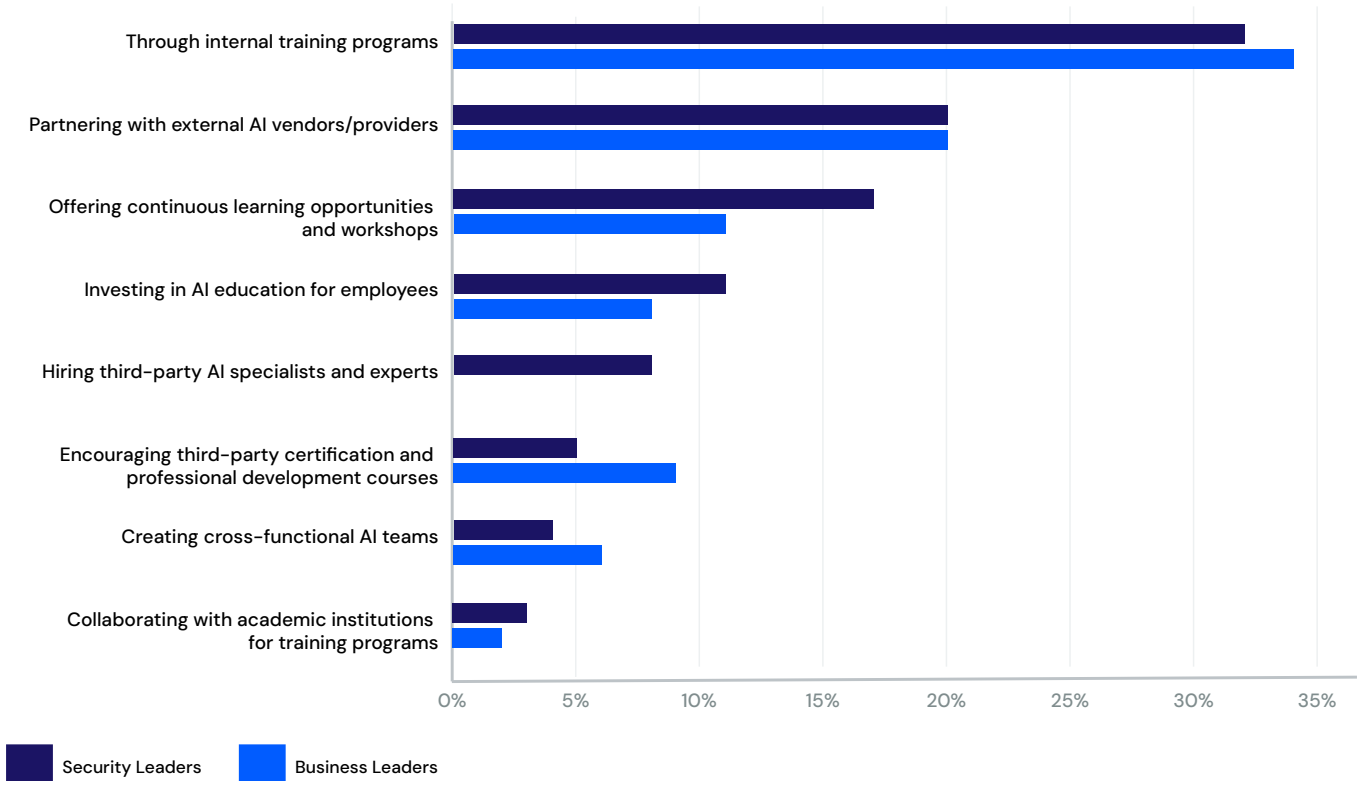

Figure 18: Addressing the Skills Gap

It's no surprise that there is a lack of skilled personnel in this highly technical niche that has just exploded across all verticals at once.

Skills gaps had earlier been identified as an issue (question 7), but there is no uniform approach to addressing this. The main route, chosen by 33% of respondents, is "through internal training programs." Next comes "partnering with external AI vendors/providers" at 20%.

Other approaches include "offering continuous learning opportunities and workshops at 14%; "hiring third-party AI specialists and experts" at 10%; "investing in AI education for employees" at 9%, "encouraging third-party certification and professional development courses" at 7%; and "creating cross-functional AI teams" at 5%, while just 2% chose "collaborating with academic institutions for training programs."

While there was general agreement between the two groups, with internal training and partnering with external AI vendors the leading routes, there were differences in opinion when it came to "offering continuous learning opportunities and workshops," which was more favored by security leaders at 17% compared to business leaders at 11%.

Other results include "hiring third-party AI specialists and experts," with security at 8% and business at 11%; "encouraging third-party certification and professional development courses" with security leaders scoring it 5% versus business leaders 9%; and "creating cross-functional AI teams," with security leaders at 4% versus business leaders at 6%.

19. How quickly are generative AI use cases being deployed from ideation to production?

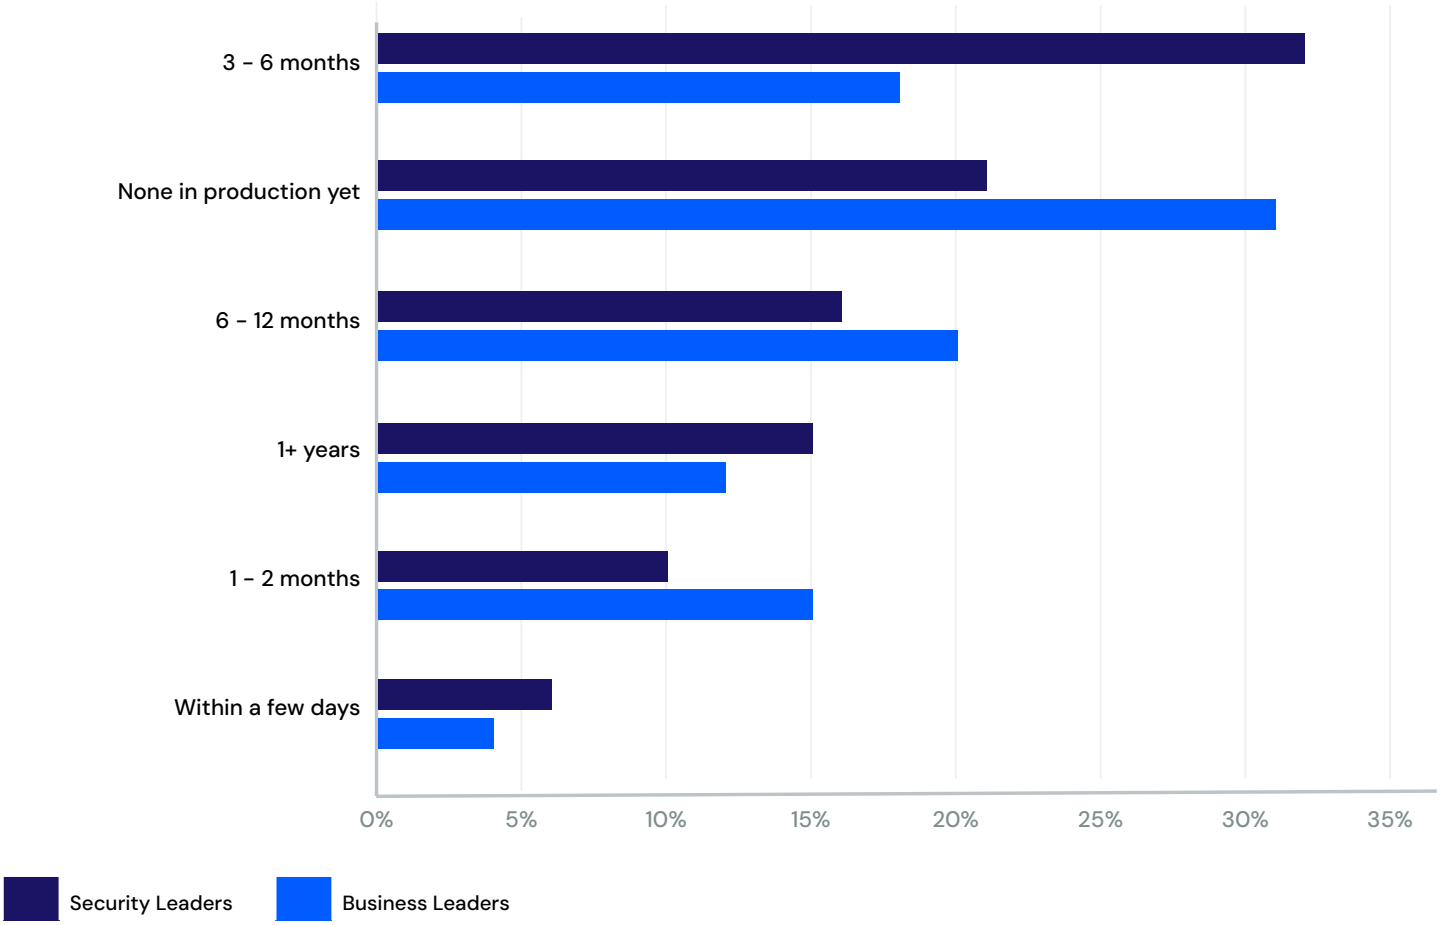

Figure 19: Perceived Speed of Deployment Varies

There are differences of opinion between security and business leaders as to how fast generative AI is being deployed - and even if it has been deployed in production.

While just over a quarter of respondents - 26% - had no AI use cases in production, of those that did, the largest group - 25% - took 3 to 6 months from ideation to production, with another 18% taking 6 to 12 months and 13% taking just 1 to 2 months.

At the fastest end, 5% report deployment within a few days, while 13% say it takes more than a year.

Perceptions differed across the two groups, while they were close when it came to those reporting a few days - with security at 6% and business at 4% - they then diverged on those reporting 1 to 2 months - with security at 10% and business at 15%. The difference became more significant when reporting 3 to 6 months, the main response for security at 32%, compared to 18% for business. The difference narrowed on 6 to 12 months - with security at 16% and business at 20%, then narrowing still further at 1 plus years - with security at 15% and business at 12%.

There is also a significant difference between those reporting none in production, with security leaders at 21% and business at 31%. Here again, there is a significant discrepancy of perception over something that should be a matter of fact, with business leaders appearing to underestimate deployment.

20. How does the pace of generative AI adoption align with leadership expectations?

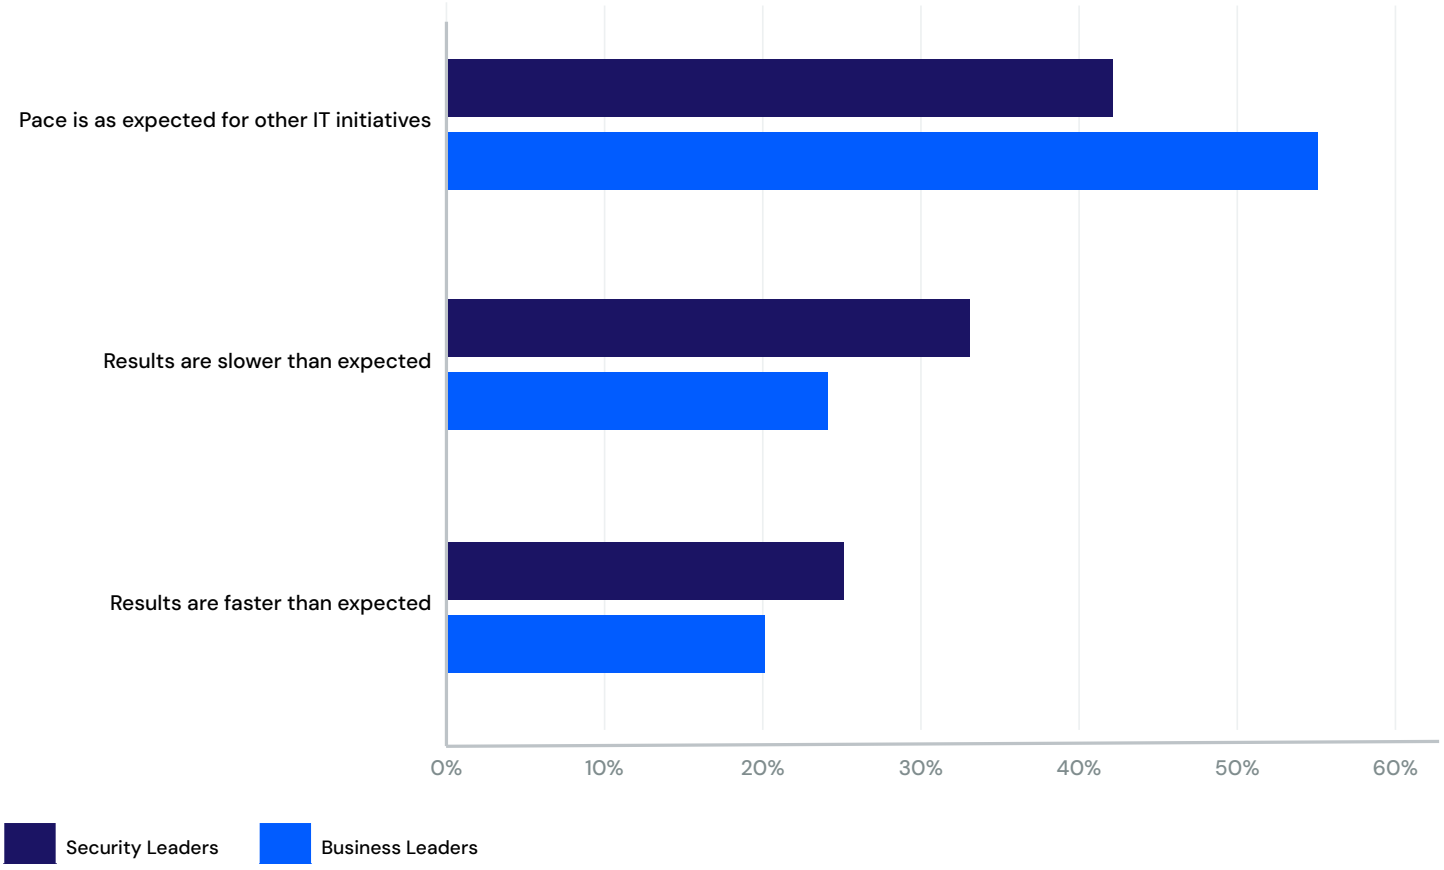

Figure 20: Generative AI Deployment on Schedule

For most organizations the deployment of generative AI is on schedule and in line with leadership expectations, suggesting ways to overcome difficulties have been found.

In broad terms, about half of respondents report deployment pace as expected - 49% - with just under a quarter saying faster than expected - 23% - and just over a quarter - 29% - saying slower than expected, indicating an even distribution.

When the figures are broken out, security leaders were more likely to report that the results were slower than expected at 33% compared to business leaders at 24%, with 42% of security leaders saying the pace is as expected versus 55% of business leaders saying the same, and 25% of security leaders saying results are faster than expected compared to 20% of business leaders.

21. Who is ultimately accountable for AI risks?

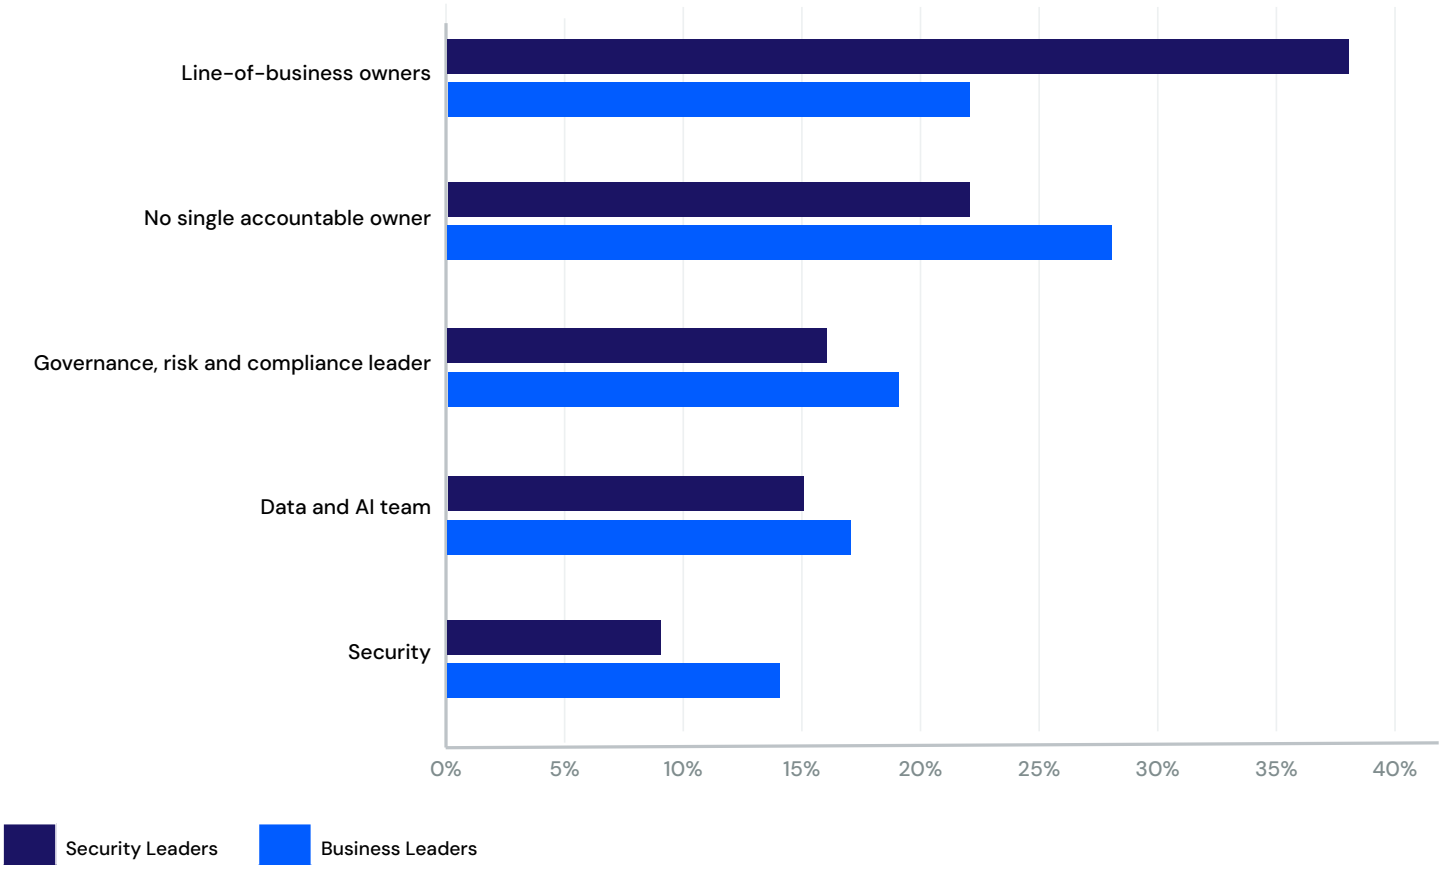

Figure 21: Generative AI - The New "Business as Normal"

While generative AI may still be seen as a largely technical issue, it is increasingly recognized that line-of-business owners will be held ultimately responsible for AI risks as it becomes an integral part of operations.

"Line-of-business owners" was the most popular response for AI risk accountability at 30%, followed by governance, "risk and compliance leader" at 18% and "data and AI team" at 16%. In a further 11%, "security" is responsible. But for 25%, there is no single accountable owner.

However, perceptions differed significantly when the two groups' responses were compared, with security leaders more likely to cite "line-of-business owners" - 38% - as ultimately accountable for AI risks compared to 22% of business leaders saying the same.

Business leaders were more likely to suggest all the other options (i.e., not one alternative), thus the divide was "data and AI team" with security at 15% and business at 17%; "governance, risk and compliance leader," with security at 16% and business at 19%; "security" with security leaders at 9% and business leaders at 14%; and "no single accountable owner," with security at 22% and business 28%.

DEMOGRAPHIC QUESTIONS

22. Job title of respondent

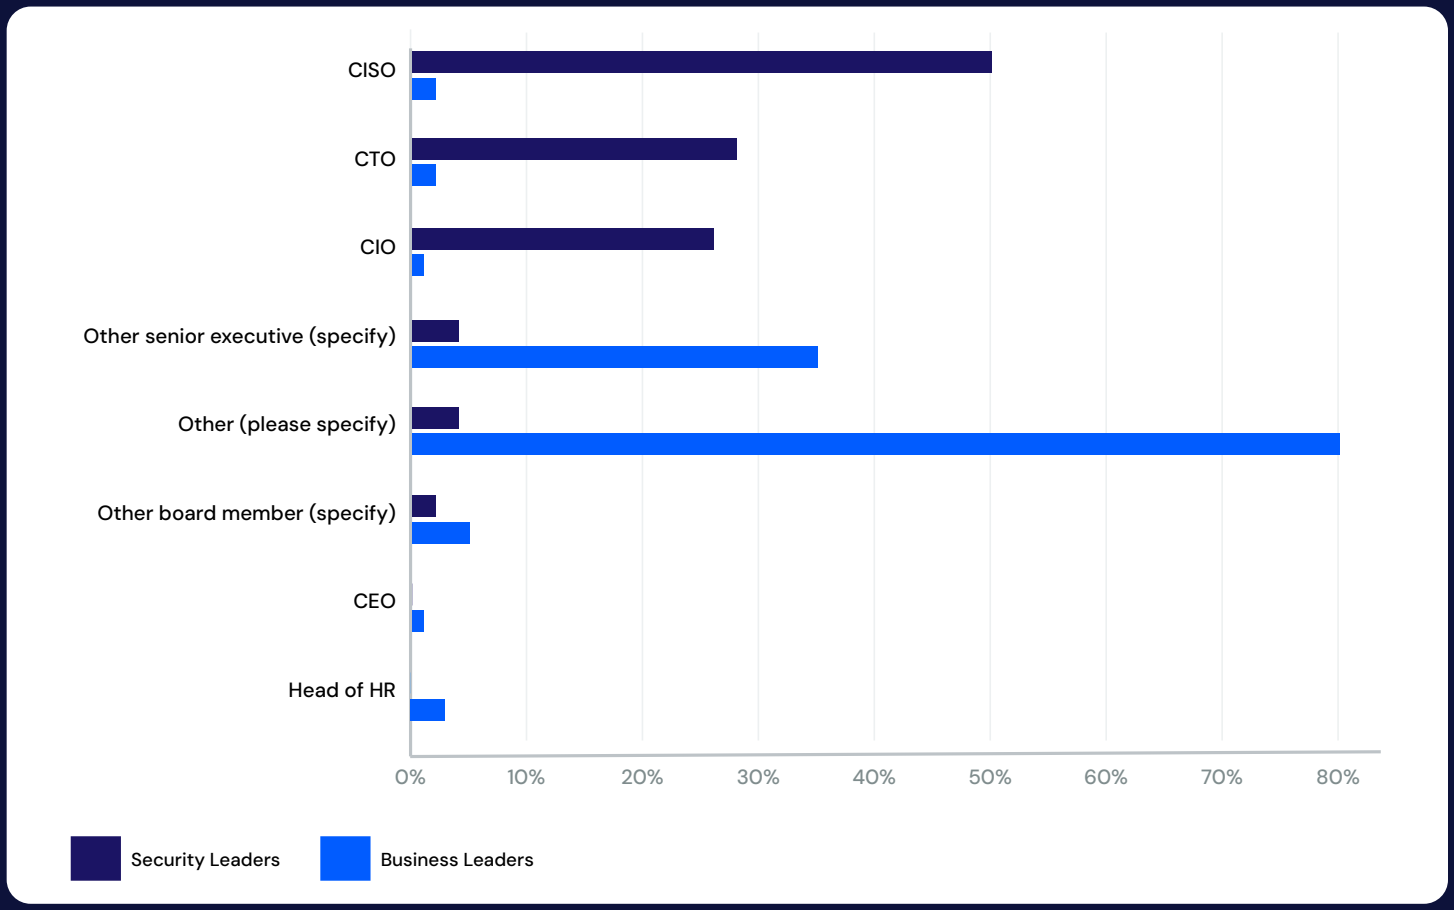

23. Respondent's industry

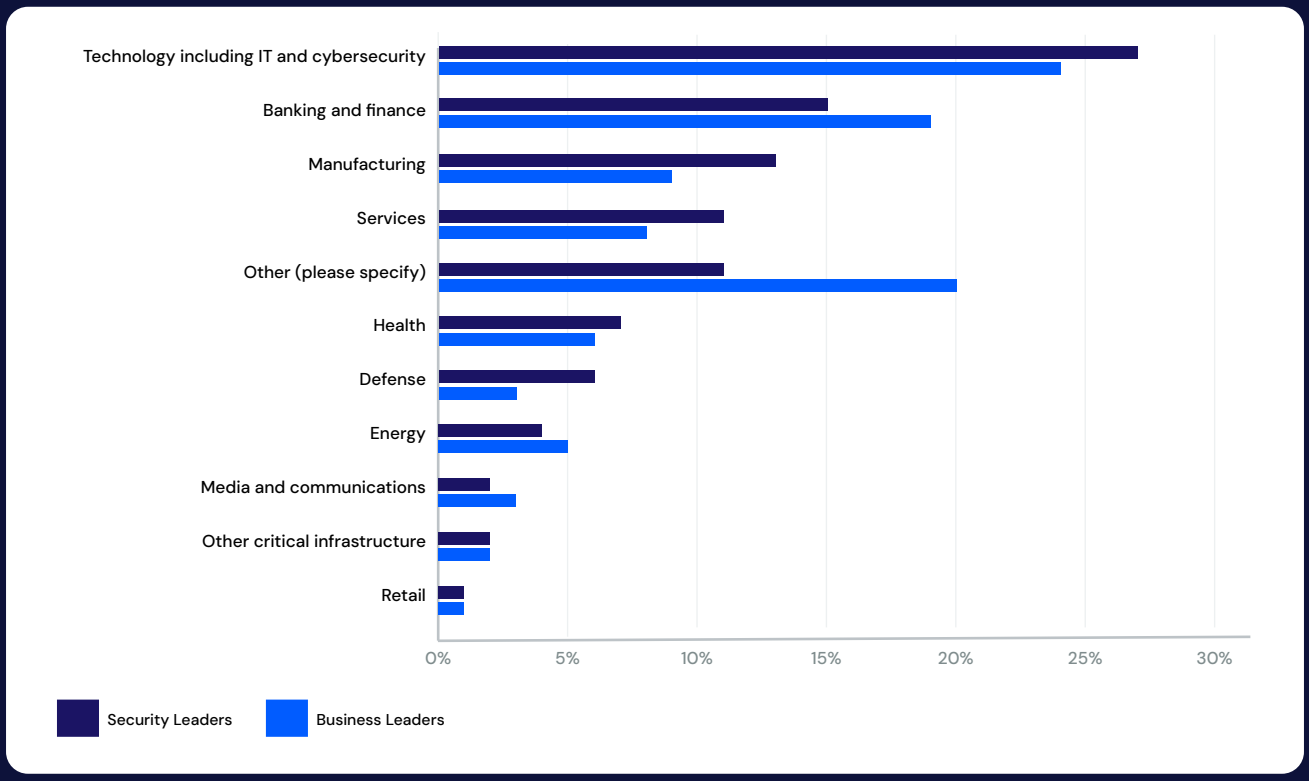

24. Respondent's geography

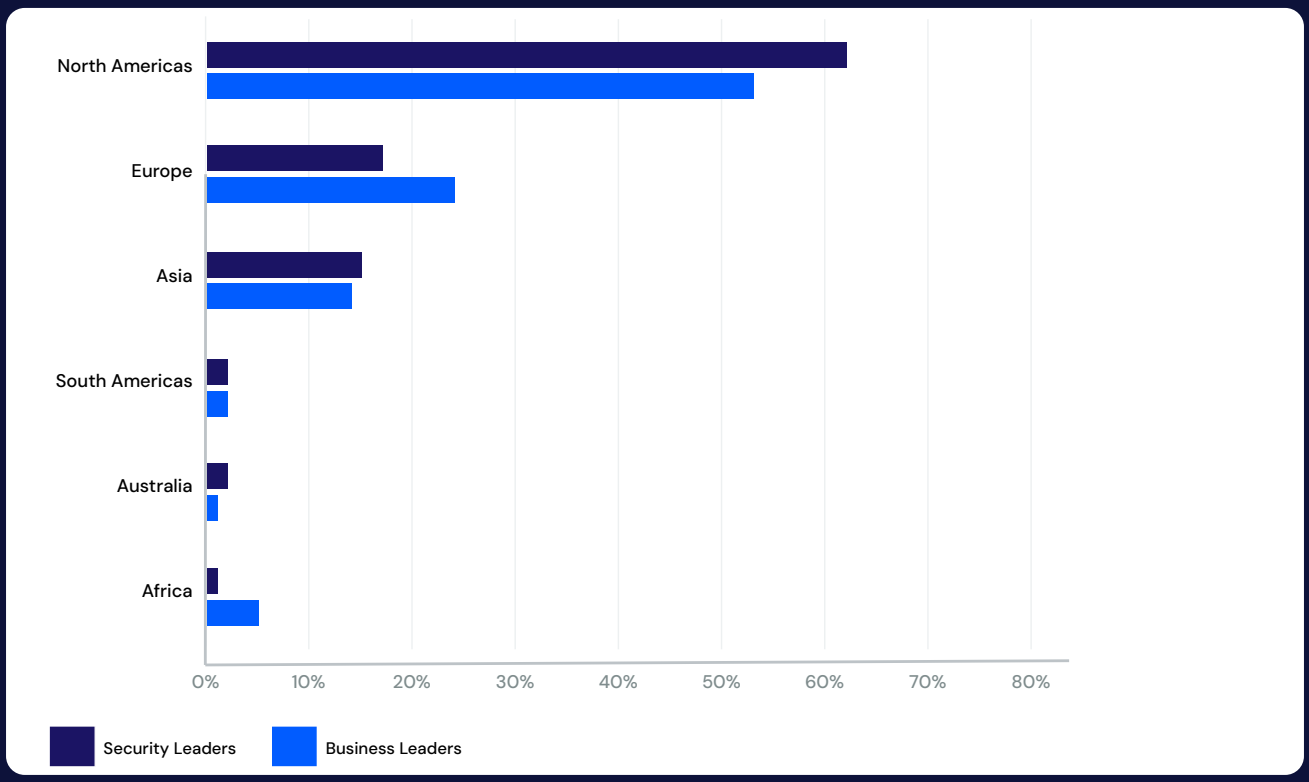

25. Size of the organization

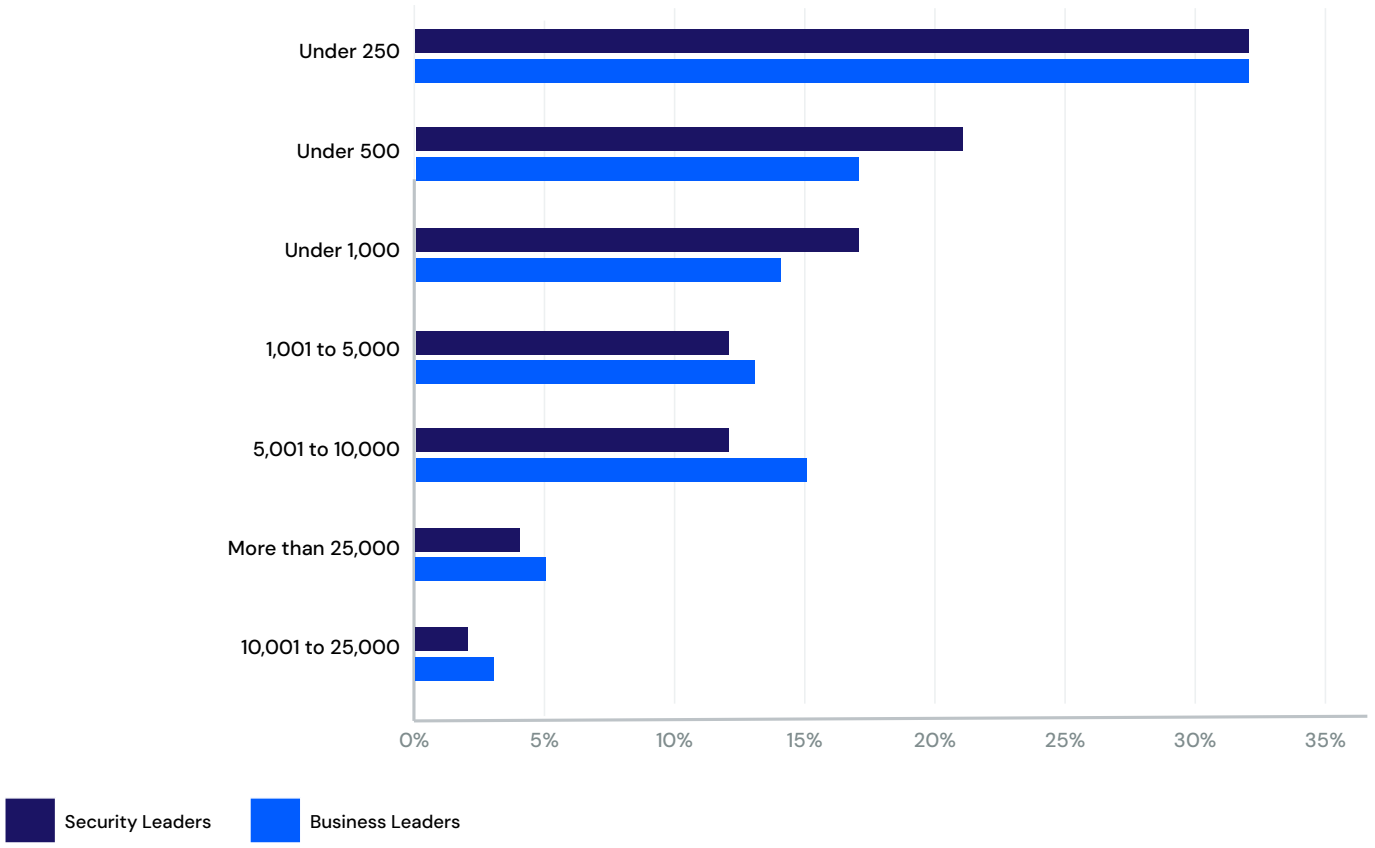

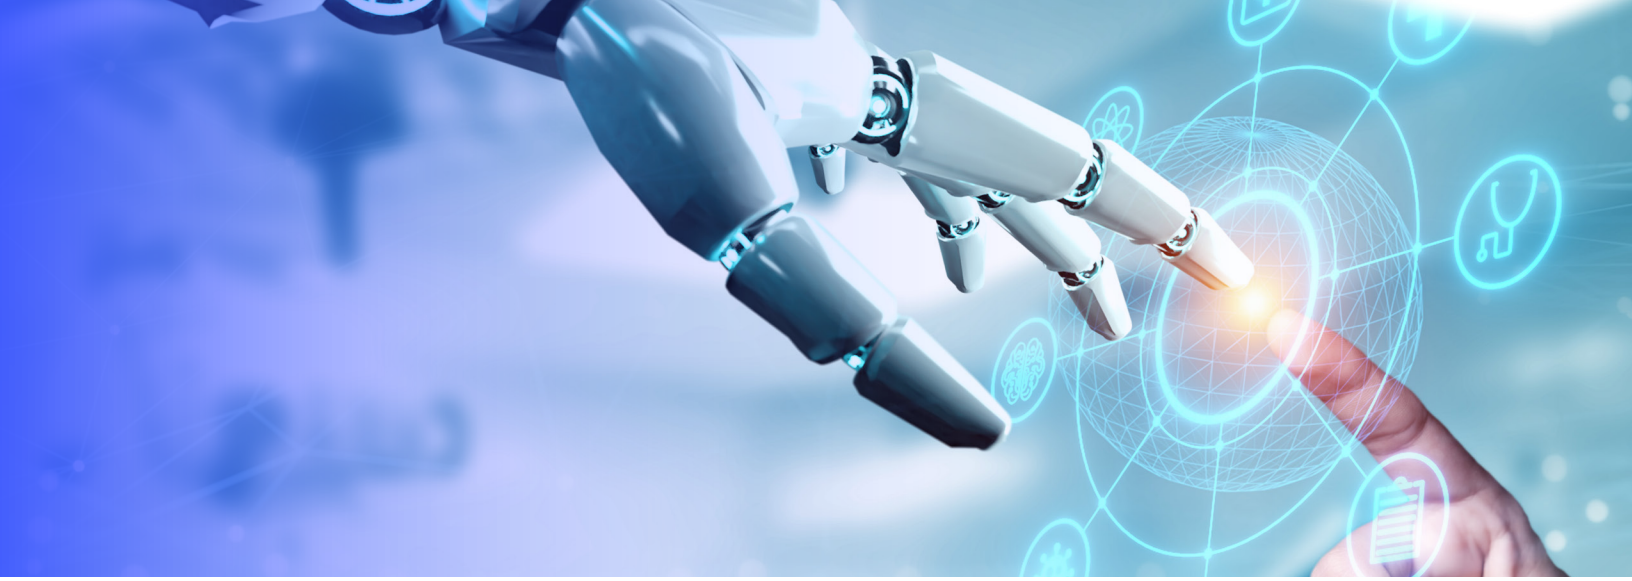

26. If you have implemented generative AI, what has been the biggest difference between what you expected and the reality? (Free text)

Free text responses varied widely, and ranged from the highly positive to the extremely negative, with the overall picture that implementation of AI is more complex and difficult than most had anticipated, with integration difficulties and lack of skilled staff among the issues encountered.

Some use cases clearly provided an excellent ROI; unfortunately, the more common response was that the potential returns of AI were difficult to achieve, if not illusory, and that the range of effective use cases is not as great as the marketing hype may have suggested.

Security responses were overwhelmingly negative, with just a few exceptions, whereas business responses, although also weighted three to one to the negative, had about one in four very positive responses.

It is likely that the outright contradictory responses reflect the diversity of use cases - with some targeting the low-hanging fruit and others being perhaps ill-suited to integrating AI, on top of which the AI models used, the existing legacy stack, the skills profile of staff involved, and the available resource (financial, staff and time to dedicate to the task) will all vary enormously.

A key takeaway is that productivity gains are achievable - in the right circumstances, but it is likely to be more difficult to implement generative AI into production securely than may have been initially envisaged, and it will not be appropriate in every case.

Responses included a range of pros and cons highlighting the differences between expectations and the reality of implementing generative AI. A representative sample of the responses is provided below:

PROS

- Huge incremental volume for our ideation and innovation practice.
- More effective than expected.
- ROI has been orders of magnitude higher than expected.
- The stability and accuracy of results, I was a bit pessimist before and once we started seeing the results, were quite impressed by its performance.
- More widely adopted than expected.
- How easy it is to experiment and find use cases.
- Time savings for some tasks.
- Repetitive tasks can be automated well.
- People's reaction was better than expected, the aversion to change in this case was not significant, probably because they quickly realized its advantages in practice.
- Improved operational efficiency.
- From a sales and marketing prospective, AI has created greater narratives and sales opportunities.
- It is very successful.

CONS

- We have trialled copilot and find that it rarely provides useful content in the "work" context. Crashes frequently and provides slow answers that often need multiple iterations of adjustment to form a workable answer. In most cases, it has been more efficient to do the work manually. If we want to implement generative AI, we will need to find a more singular/focussed platform (e.g., chatbot or email response composing mechanism) that is able to feed in accurate context from our sources of knowledge.
- Time to get the hardware working has been slower than anticipated and far more expensive.
- We now know, too late, that the products couldn't possibly deliver the promised/promoted results. They simply are not ready for prime time despite all the marketing hullaballoo.
- Senior management and stakeholders do not know what generative AI does for them as most of them do not know how to apply it for their ROI.
- Learning how to write prompts to get the desired output is more challenging for our users than I thought it would.
- Integration with existing applications has been difficult.
- It is not mature yet and requires lot of effort for implementation.
- Bias in training materials significantly impacts AI predictive analytics.
- Requires extensive training and testing before deployment.
- Expectations that media portrays as the solution for all use cases/problems but the reality is that only certain use cases are feasible at present.

CONS

- The biggest difference has been the complexity of integrating AI into existing cybersecurity frameworks. While AI tools were expected to automate security tasks, managing their deployment and ensuring they work seamlessly with legacy systems required more time and expertise than anticipated. Additionally, the management of AI-driven outputs requires careful monitoring to prevent unintended security vulnerabilities.
- Takes longer to implement, more complex.
- How little employees know how to use it; how little leadership values its potential.
- We are still learning new ways to prevent hallucinations.
- The reality is that to make generative AI tools secure and compliant with fast-moving laws and technology is a challenge in itself. In reality, we see that AI has a hallucination rate and what it does today is not all as stated in the vast majority of marketing articles. We need to engage and do hands-on to identify the reality, then translate this without marketing part in a practical way to business who then can build use cases based on reality.
- The squeeze isn't worth the juice.
- It does not have good ROI in all scenarios and applications.
- It is yet to be seen where we are seeing material gains from implementing AI.
- Does not live up to the hype.
- It is hard to implement as we lack the skills we thought we had.

CONCLUSIONS

GENERATIVE AI DEPLOYMENT

Implementation of generative AI use cases remains a tech-driven process, with a large security element; however, the extent to which line-of-business management is taking responsibility continues to increase, suggesting a rapid move from the tech domain toward "business as usual" for generative AI use.

Currently, there is often a failure of communication resulting in significant discrepancies between how security and business leaders understand and view the deployment of generative AI in production, which, in some cases, exacerbates differences in prioritization, beyond what might be expected from the two roles' different objectives.

While there are differences in perception, it is clear that actual deployment of generative AI in production has risen significantly, with averaged responses showing a rise from 15% to 31% over the past year, and when AI pilots are included, well over half of respondents - 63% - now have some form of generative AI deployment.

Automation of repetitive tasks leads the priority list of use cases, cited by 82% of respondents, up from a year ago, with the second-placed objective to "increase speed of production/ service" also up at 67% and third-placed "perform routine and administrative tasks" at 62%, suggesting increased focus on these core areas as achieving the best ROI for early implementations.

Office functions such as customer service including chatbots, sales, and marketing, and also tech applications such as security and software development, were important domains for generative AI use, though understandably security leaders rated tech deployments higher than business leaders.

GENERATIVE AI CONCERNS

A year ago, fears around the security of generative AI led nearly a third to introduce bans, saying they had not nor did they intend to implement generative AI; only 6% now say the same.

While the levels of concern about generative AI implementation have fallen, they do remain high, led by "leakage of sensitive data by staff using AI" cited as a top concern by 76% of respondents - down from 81% a year ago - with potential impacts on reputation, compliance and market competitiveness. "Ingress of inaccurate data (hallucinations)" remains high at 59%, but has also fallen from 69% a year ago, while ethical concerns remain consistently high.

When it comes to challenges faced when integrating generative AI with existing IT infrastructure, ensuring data privacy and security remains top at 73%, but while high, it is

down from 80% a year ago. The lack of skilled personnel for implementation remains significant at 60%, along with other challenges such as data integration and migration difficulties, the high costs of deployment and maintenance, and compatibility issues with legacy systems.

Clearly, concerns and challenges persist, but tools and approaches developed to address them seem to be making an impact in reducing their intensity.

When concerns and challenges are viewed in conjunction with the deployment figures, it appears that in assessing the risk/return equation, more respondents now believe they cannot be left behind, and that the risk of inaction is deemed greater than the risk of deployment.

ROI EXPECTATIONS

Notwithstanding security concerns, both groups are looking to generative AI to achieve operational efficiency gains and more than half of respondents now have specific generative AI-related purchase plans, with over a quarter now having a specific AI budget - which is twice as much as a year ago.

Forecasting of productivity gains was generally quite optimistic, despite the problems faced, with 67% of respondents predicting 6% to 30% productivity gains.

When comparing with previous figures, it is difficult to judge whether the early adopters in last year's report were achieving gains based on targeting the low-hanging fruit and are now seeing less substantial gains, or that more respondents are in a position to judge based on more robust use cases and maturing generative AI models. Business leaders are generally more optimistic about the gains, but also represent both more pessimists and more optimists, while security leaders are more grouped in the moderate optimistic range.

While there is optimism about the benefits and an increase in the expected ROI on investing in generative AI, our free text responses show that the actual returns achieved are hugely variable, from those vastly exceeding expectations to a majority reporting disappointment with returns.

Key factors impacting underperformance include an underestimation of the complexity

and associated difficulties involved in integrating generative AI use cases and technology into legacy systems. This included finding out during implementation that staff skill sets were not appropriate to the tasks being asked of them. It also appears that there is huge variability in returns depending on the use cases for which generative AI is deployed. This is to be expected as some low-hanging fruit will be well-suited for automation with generative AI, while in some cases, either the company or the AI model is not appropriate for the task.

UNDERSTANDING IS UP

Even though there has been a rapid change with an increase in the number and variety of regulations and the amount of guidance globally that impacts the AI market, there is a reported increase in understanding regarding what restrictions apply, up from 45% to 61%. This indicates that the importance of both AI itself and the need to progress in a compliant manner is widely understood - though the extent of those not understanding the regulations that apply, at 40%, remains very high for such an important aspect of doing business going forward, acting as a reminder that widespread generative AI use is a nascent industry.

APPROACHES TO COUNTER CONCERNS

The leading approaches to tackle security are led by staff education and training around the secure use of AI, which garnered more than 60% of responses as it becomes increasingly accepted that whatever technology is implemented, adversaries will target human to sidestep the technology. That said, technological approaches such as encryption of data are also supported by more than 52% and process/rule-based approaches are widespread.

In an earlier question, only 6% said they did not use or plan to use AI, yet this is contradicted when we specifically looked at ways companies were overcoming concerns, with 14% now saying they would ban all generative AI use. Banning use of all generative AI was reported by 11% of security leaders and 16% of business leaders.

Bans on personas/departments are also supported along with use of blocking software to prevent both the ingress and export of unauthorized data.

It was interesting that the technologists, the security professionals, were significantly more likely to cite staff education and training around the secure use of AI as the primary tool to mitigate their concerns around AI use, compared to business leaders who favored technological solutions. That could be because the security leaders are more aware of the limitations of technology, or equally it could be because the business leaders have taken a data-driven ROI approach.

Defining those gains, improvement in operational efficiency was the highest-ranked ROI measure at 74%, closely followed by cost savings from automation at 67% and time saved on repetitive tasks at 65%. In addition to these bottom-line targets, more than half of respondents - 53% - cited better customer satisfaction and engagement.

There has certainly been a change over the past year, with increased optimism about the longer-term gains from generative AI, an increase in deployment, and a reduction in the still high level of security concerns as tools to mitigate those concerns roll out. However, there is also a realization that the practical rollout and integration of generative AI use cases into existing infrastructure is complex, difficult, and can eliminate gains if not managed well by appropriately skilled staff - who are in short supply.

EXPERT ANALYSIS

Christine Livingston, Managing Director, Global AI Leader, Protiviti

Livingston has over 18 years of experience in technology consulting and applies her engineering background with advances in technology to unlock enterprise innovation. With over a decade of experience in AI-ML deployment, she has delivered hundreds of successful AI solutions, including numerous first-in-class AI-enabled applications. She has helped many Fortune 500 clients develop practical strategies for enterprise adoption of AI.

GENERATIVE AI RESISTANCE GIVES WAY

MORBIN: What's driving the sharp decline in organizations prohibiting generative AI use - reduced concerns about generative AI, greater confidence about the technical and procedural protections, fear of missing out on the potential gains, or something else like shadow AI?

CHRISTINE LIVINGSTON: There are a couple of reasons that we see that dropping. First is probably that even among organizations that outright blocked a lot of the public generative AI sites, they found that a large percentage of their employees were still leveraging those technologies. We've seen that consistently across many different organizations. So a pure block or a ban was not overly successful. They've also started to see some of the early-stage pilots and prototypes come to fruition in production. The number of projects in production almost doubled in the last year.

But as you start to see those use cases move into production, you begin to recognize the tangible value that AI is creating for your organization, and it becomes much harder to ignore the potential and the possibilities. We have also come a long way in terms of understanding some of the potential pitfalls and risks of generative AI specifically and have better opportunities to mitigate and manage those potential downfalls.

GROWTH TRAJECTORY OF AI ADOPTION

MORBIN: With general AI usage in production doubling from 15% to 31% in a year and given all the hype around AI, is the growth slower than you would have expected or is it doubling faster?

LIVINGSTON: It mirrors generally what we're seeing. It's maybe a little bit slower than I would've expected, but at the same time, we know that a lot of these applications are first in class. They're proving potentially an unproven

concept applied to a new business process, and that requires a lot of planning and a lot of technical validation that the solution will work the way you expect it to. So if you look at the results from last year, about 27% of the respondents last year said they had plans to implement AI. About 28% of people said they were in the pilot phase last year.

So if you take those 28% in pilot last year and think about where they are a year later, you're almost exactly at that growth trajectory that we would've expected. So as people are budgeting, experimenting and then releasing to prototype in POC, this naturally becomes the trajectory of the solution into production.

UNDERSTANDING AI ROI VARIATIONS

MORBIN: Productivity gains seem to vary widely, with some reporting exceptional ROI while others consider any ROI an illusion. What factors drive this disparity in responses?

LIVINGSTON: One of the most significant challenges faced by organizations today is selecting the appropriate use case. Many organizations have chosen what appeared to be the simplest option, without fully understanding the business impact or value of utilizing AI in that function. They may not have quantified an ROI before they began experimenting. I've seen a huge range of value delivered based on the use case that clients look to pursue and the organization's approach.

The other concept tied to ROI is how most people who are experimenting with this technology use it as a consumer. They're familiar with how the technology works. However, it's much harder to operationalize an enterprise use case for generative AI than do some light experimentation. Enterprises often underestimate the challenges of integrating AI with their data, establishing governance principles, setting up guardrails, and ensuring a justifiable business case. These complexities ultimately influence ROI outcomes.

One of the most significant challenges faced by organizations today is selecting the appropriate use case. Many organizations have chosen what appeared to be the simplest option, without fully understanding the business impact or value of utilizing AI in that function.

– Christine Livingston, Managing Director, Global AI Leader, Protiviti

KEY AI RISK CONCERNS

MORBIN: Are the concerns shared by respondents, particularly around data leakage and unreliable results, aligned with what you or your organization identifies as the most important risks or do you see other risks as more critical?

LIVINGSTON: Unreliable results, often referred to as hallucinations, are probably one of the predominant concerns for organizations today. How do we trust the answers? There are also a lot of concerns about how we retrain our employees and our organizations to use these results appropriately and accurately. I often use the analogy of reviewing a spreadsheet or a forecast: I know which cells to look in and which formulas to validate to confirm the accuracy of the data I'm seeing and using.

With generative AI, we don't necessarily know yet where to look and how to validate the responses and outputs. Absent a clear citation of source data, we're still learning those behaviors and habits around how to interpret and use the responses and the outcomes responsibly. A lot of organizations are moving these capabilities in-house - within the confines of their cloud platforms - to ensure their data is not being used to retrain models or entering the public domain. These, in my view, are probably the top risks that clients are most concerned about today.

MITIGATING GENERATIVE AI RISKS

MORBIN: What should we do to mitigate the risks associated with generative AI?

LIVINGSTON: One key point I emphasize is that not all risks are necessarily created equal. Organizations should focus on stratifying risks into low, medium and high categories and designing governance processes, policies and technologies accordingly. These discussions are complex and need to be centered on the use case to make sure that your risk mitigation is appropriate.

Retrieval-augmented generation (RAG), output comparison and model evaluation, are tactics that are often used. However, a significant challenge for many organizations lies in performing the initial risk classification and turning it into actionable strategies. Developing a target operating and governance model to support from the people, process and technology standpoint is essential for effective risk management.

AI IN GOVERNANCE CHALLENGES

MORBIN: Why is AI adoption in governance limited - is it due to a lack of understanding or caution in using it for critical services?

LIVINGSTON: The response reflects the natural tension we often see in organizations today. On one side, there are innovation-minded individuals who are eager to move quickly and deploy new technology. On the other side, there are some skeptics and risk-minded colleagues advocating for a cautious, responsible approach to a thorough understanding before proceeding.

It's not surprising that the same group that's responsible for governance and risk management is probably not the first adopter of the technology. They're likely in a "wait-and-see" mode, focusing on ensuring that some of the other use cases they're responsible for are well-managed, and they're confident with some of the risk mitigators in place.

Additionally, there is still an emerging level of understanding of how AI can be used in this function. We have seen interesting applications of AI in governance functions, such as identifying overlooked risks, evaluating 54 whether controls are sufficient or accurate, and highlighting potential considerations. But generally, I would say the response and findings are not overly surprising to me.

NOTABLE SURVEY INSIGHTS

MORBIN: Are there any other results that stood out or surprised you, or anything you'd like to add that we haven't covered?

LIVINGSTON: The breakdown was particularly interesting. What stood out to me was the year-over-year shift: a year ago, more business leaders reported using generative AI than security leaders, and now that trend has reversed. This change likely ties into the conversation around governance and the operationalization of use cases. As organizations move toward pushing solutions into production, security teams become more involved, gaining greater visibility into AI usage. The reversal is a fascinating development and seems to reflect the progression of AI solutions in production.

ABOUT THE SPONSOR

protiviti®
Global Business Consulting

Protiviti (www.protiviti.com) is a global consulting firm that delivers deep expertise, objective insights, a tailored approach and unparalleled collaboration to help leaders confidently face the future. Protiviti and its independent and locally owned member firms provide clients with consulting and managed solutions in finance, technology, operations, data, digital, legal, HR, risk and internal audit through a network of more than 90 offices in over 25 countries.

Named to the Fortune 100 Best Companies to Work For® list for the 10th consecutive year, Protiviti has served more than 80 percent of Fortune 100 and nearly 80 percent of Fortune 500 companies. The firm also works with government agencies and smaller, growing companies, including those looking to go public. Protiviti is a wholly owned subsidiary of Robert Half Inc. (NYSE: RHI).

About ISMG

ISMG is the world's largest media organization devoted solely to cybersecurity and risk management. Each of its 38 media properties provides education, research, and news that is specifically tailored to key vertical sectors including banking, healthcare, and the public sector; geographies from North America to Southeast Asia; and topics such as data breach prevention, cyber risk assessment, AI, OT, and fraud. Its annual global summit series connects senior security professionals with industry thought leaders to find actionable solutions for pressing cybersecurity challenges.

Contact

(800) 944-0401
info@ismg.io

Sales & Marketing

North America: +1-609-356-1499
APAC: +91-22-7101 1500
EMEA: + 44 (0) 203 769 5562 x 216

CAREERS INFO SECURITY · BANK INFO SECURITY · GOV INFO SECURITY · HEALTHCARE INFO SECURITY
CU Just for Credit Unions INFO SECURITY · Data Breach Today · InfoRisk Today · AIToday.io · OT.today · CIO.inc
CyberEd.io · CyberEdBoard · DeviceSecurity.io · FraudToday.io · PaymentSecurity.io
CYBER THEORY · GREYHEAD AN ISMG COMPANY · Xtra mile LIFECYCLE MARKETING · QC MEDIA · ATHENA

ISMG

902 Carnegie Center • Princeton, NJ • 08540 • www.ismg.io