

Google Cloud

exabeam

Clearwater

onetrust

Microsoft Security

FIRST ANNUAL GENERATIVE AI STUDY:

Business Rewards vs. Security Risks

TABLE OF CONTENTS

Welcome Summary	3
Introduction.....	4
By the Numbers	6
Executive Summary	7
Survey Results.....	9
Expert Analysis	29
Conclusions	41
About the Sponsors.....	42

Summarizing the First Annual Generative AI Study Report

This Q3 2023 global survey involved over 400 anonymous professionals in business and cybersecurity, sharing insights from two distinct groups: business leaders - 47% - such as CIOs, CEOs, and executives, and cybersecurity professionals - 53% - predominantly at the CISO level, with varied titles including head of security, head of IT and head of SOC. The study provides a comprehensive snapshot of perspectives within these key roles in sectors including finance, manufacturing, healthcare, tech/biotech, government, critical infrastructure and retail, among others.

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

This survey of over 400 business and cybersecurity professionals conducted in Q3 2023 comprises responses from two cohorts, business leaders – comprising 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, and the largest group comes from North America.

In the survey, we look at the differences in perspective between business leaders and cybersecurity professionals when it comes to their current and intended use cases for generative AI. Where generative AI is deployed, we look at measuring productivity gains, and where it is not currently used, we look at the anticipated gains and intended deployment. This includes current and intended allocation of expenditure and its projected growth as well as areas for investment going forward.



INTRODUCTION

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

Also, the survey seeks to get a snapshot of current understanding of generative AI, including the range of generative AI tools being explored/trialed, as well as respondents' understanding of current regulation.

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 introduction of generative AI. This report benchmarks what your competitors are doing so that you can use these results to help enhance your own defenses and identify the productivity opportunities that gen AI presents.

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



Tony Morbin
Executive News Editor, EU
Information Security Media Group

Morbin is a veteran cybersecurity and tech journalist, editor, publisher and presenter working exclusively in cybersecurity for the past decade – at ISMG, SC Magazine and IT Sec Guru. He previously covered computing, finance, risk, electronic payments, telecoms, broadband and computing, including at the Financial Times. Morbin spent seven years as an editor in the Middle East and worked on ventures covering Hong Kong and Ukraine.

Statistics that jump out from the First Annual Generative AI Study: Business Rewards vs. Security Risks:

15%

of respondents currently implement gen AI.

13%

of respondents have a specific budget for gen AI solutions.

48%

of cybersecurity leaders do not understand AI regulations that apply to their sector.

62%

of business leaders do not understand AI regulations that apply to their sector.

When comparing the responses of business leaders and cybersecurity professionals in relation to their views on implementation of generative AI, this report finds that business leaders – while aware of the risks – are generally more enthusiastic about adopting generative AI than their cybersecurity compatriots. They are more likely to report using or trialing gen AI, and they are doing so via a wider variety of AI iterations. They are also less likely to say that gen AI has no place in their operation.

AI Deployment and Productivity

In contrast, cybersecurity professionals – while aware of the productivity opportunities for deployment in their own sector – have a higher level of concern about the risks entailed and how they might be mitigated.

Among all respondents, there is roughly a 70/30 split between those keen to adopt AI and those currently rejecting its use or who are in organizations/roles where its use is not allowed. Outright bans on use of generative AI are reported more frequently among cybersecurity professionals than business leaders, but it is not an uncommon response to tackling the risk.

More than half of all respondents who say they are actually deploying AI report more than 10% productivity gains, and some report substantially more. At the lower end of productivity gain, twice as many cybersecurity professionals - 27% - report gains of less than 5%, compared to business leaders at 14%.

For both business leaders and cybersecurity professionals, 13% report having a specific budget for generative AI thus it is clearly still at an early stage in enterprise rollout and budget cycles.

Understanding of AI

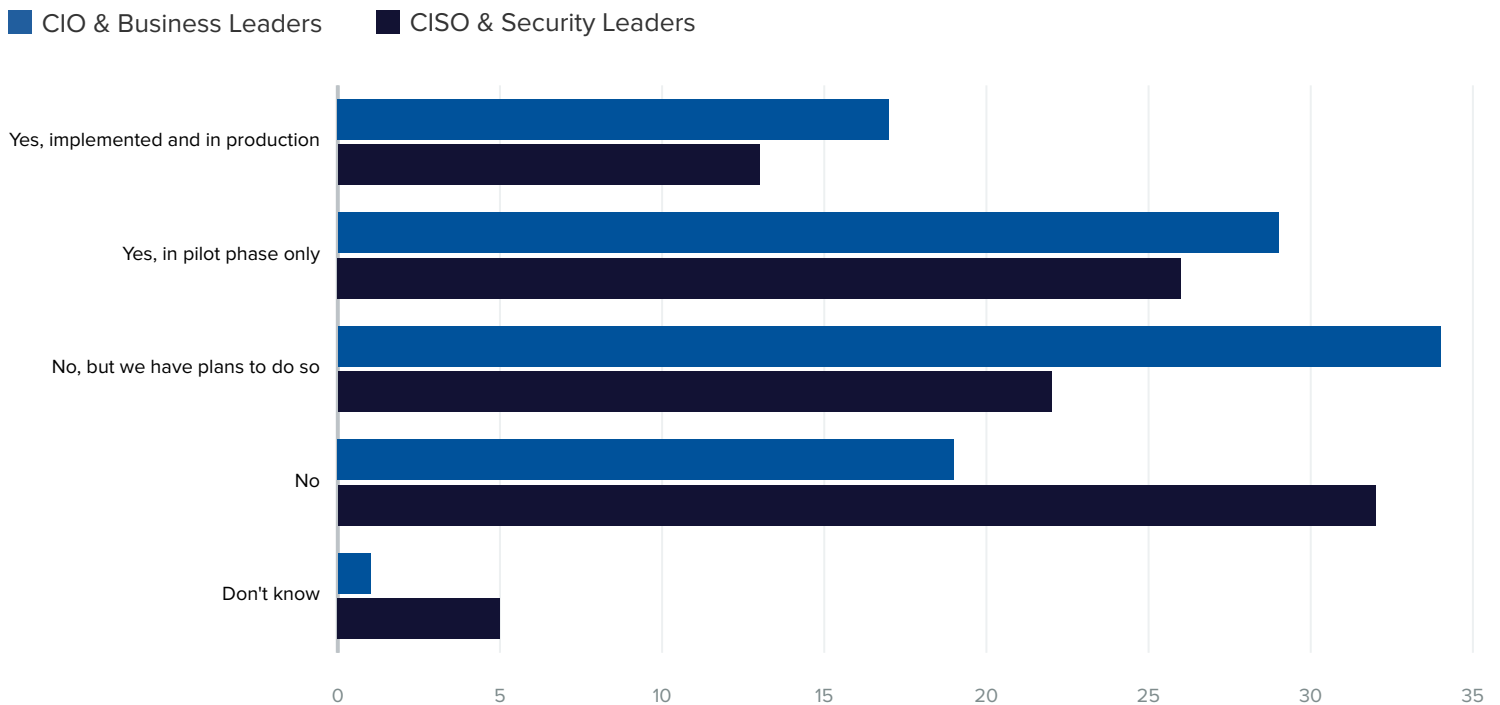
The top concerns about use of AI are leakage of sensitive data by staff using AI, cited by 80% of business leaders and 82% of cybersecurity professionals. Second for both groups is ingress of inaccurate data - hallucinations, which is cited by 71% of business leaders and 67% of cybersecurity professionals.

Particularly significant is that 38% of business leaders and 48% of cybersecurity leaders expect a continued ban on all use of generative AI in the workplace. Also, 73% of business leaders and 78% of cybersecurity professionals say they intend to take a walled garden/own AI approach going forward.

Regarding understanding of AI regulations, a worryingly low 38% of business leaders say they do understand these regulations, as do 52% of cybersecurity leaders. Yet these figures should not be surprising given the rate of change and lack of universally accepted standards and regulations.

Throughout the survey, more cybersecurity professionals than business leaders give the answer “Don’t know,” which is unsurprising since business leaders would be more expected to know their organization’s plans.

1. Does your company currently use generative AI?

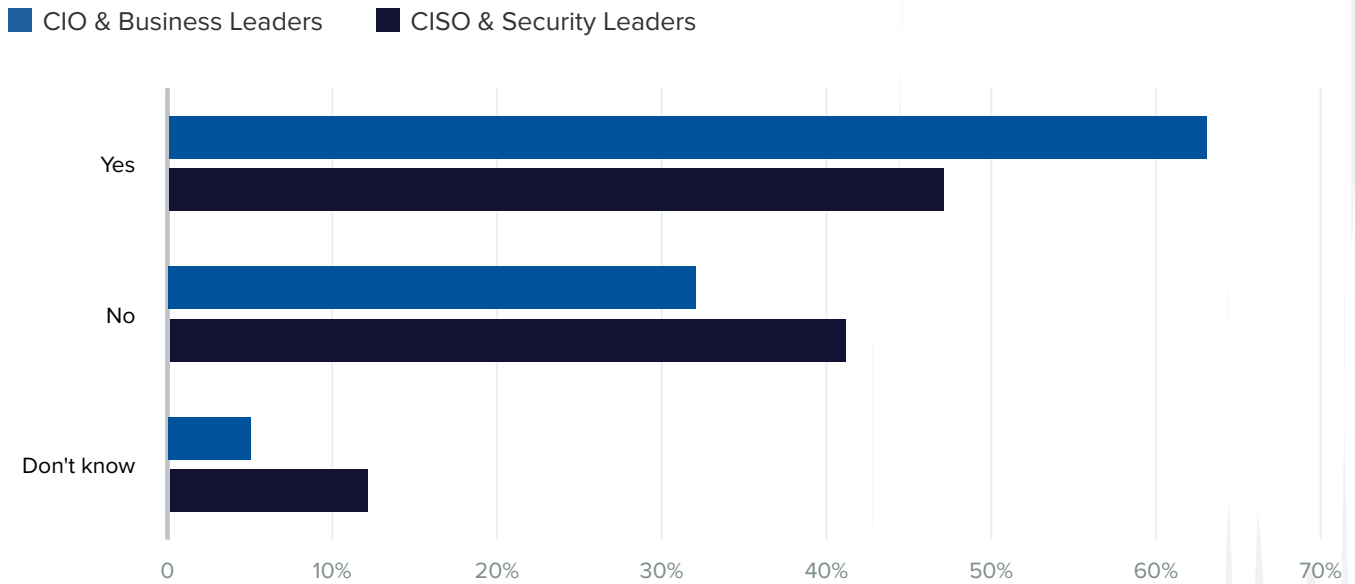


While 15% of respondents say they have deployed generative AI in production, another 28% are piloting it, bringing total current usage to 42%.

There are also 27% of respondents who neither use it nor plan to, potentially reaching 30% if we add the 3% who say they are “unsure.”

The business leaders are between 5% and 10% more likely to report implementation of AI than cybersecurity professionals, however the trend flips when examining those with no planned implementation. Cybersecurity professionals are significantly more likely to have no plans, at 34% compared to just 19% for business leaders.

2. Does your organization allow staff to use generative AI for work purposes on their own initiative?



Business leaders take the lead, with 63% reporting that generative AI use is allowed, compared to 47% of cybersecurity professionals. Within the “not allowed” camp, cybersecurity professionals hold a slim majority with 41% versus 32%.

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

The most frequent answer from business leaders is CIO. CTO and CEO are also mentioned. Other titles mentioned included IT, COO and various heads of projects/products – plus the poignant “nobody” and more enigmatic “Still a bit of a mystery.”

Most cybersecurity professionals answer CTO, with CIO not far behind, followed by CEO and CISO. Other responses include “not allowed,” “not decided” “no one” and “don’t know.”

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

The leading title mentioned by business leaders was CISO/CSO, followed by CIO, CTO and CEO or president at 6%. Cybersecurity professionals answer CTO, with CIO not far behind, followed by CEO. IT gets several mentions.

5. Who in your organization will be responsible for ongoing management of generative AI productivity solutions (job title)?

Among business leaders, the answers are led by CIO, followed by CTO. There are just a few CEOs and IT departments mentioned, and a lot more say “don’t know” or “undecided.”

For cybersecurity leaders, the CIO and CTO have roughly equal representation, with even more answering “don’t know”, “undecided” or “We’re figuring it out in the pilot.”

It is likely that the business leaders look to board peers or next level down in seniority to take an AI lead, and having a broader view than the security aspect, many will see it as primarily an information governance issue and will cite CIO. The smaller group of business leaders who see AI as a technical issue will cite CTO. In contrast, cybersecurity professionals are likely to be more focused on the security and

technical aspects and thus give a higher weighting to CTO than the business leaders.

In both cases, the high number of “don’t knows” reflects the infancy of AI’s introduction into the business sphere for most organizations. Cybersecurity professionals are less likely to make the decision to allocate authority and thus less likely to know where responsibility will lie.

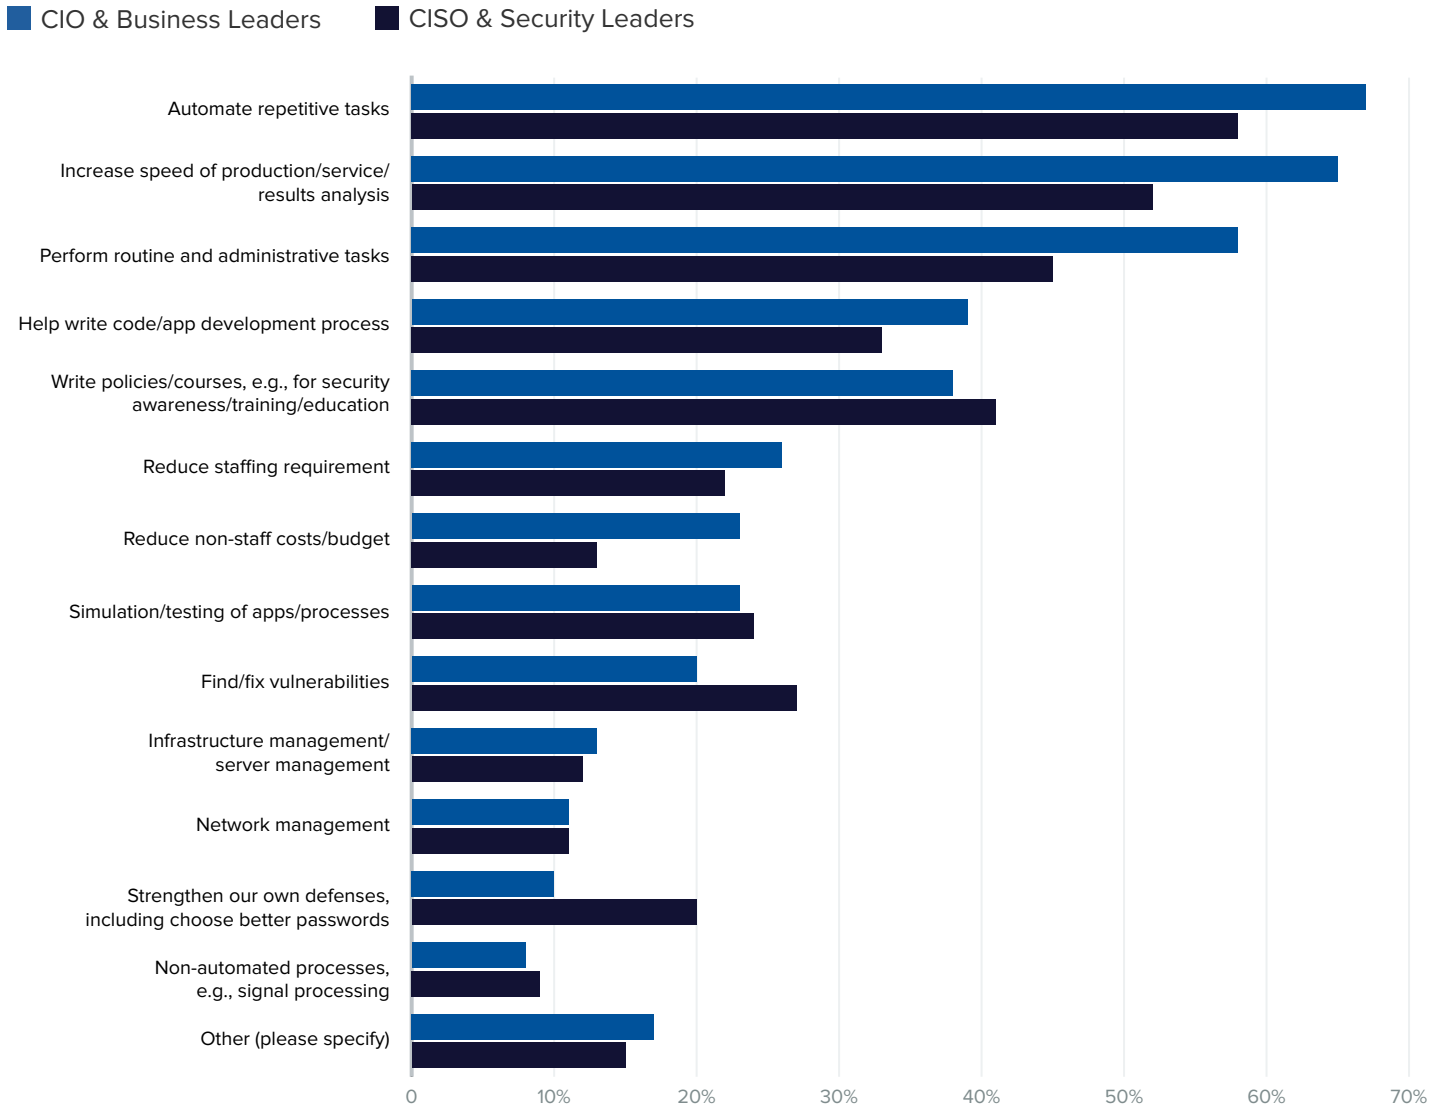
6. Which of the generative AI tools/platforms do you use, or are aware of?

Both groups say Chat GPT/GPT4, followed by Google Bard and Bing. Midjourney was also often mentioned by business leaders. It appears that they are experimenting a lot with new gen AI entrants as each offering scrambles to establish itself as a niche leader, looking to see how they can grasp the productivity gains that might be delivered.

Cybersecurity leaders also most frequently cite Chat GPT, Google Bard and Bing, and other providers are rarely mentioned. A likely explanation is that generative AI is not yet robust enough for many critical cybersecurity applications, and the operational nature of cybersecurity tasks demands more proven and tested solutions.

Also, many respondents say that no generative AI is currently used as generative AI tools are not approved for use or not allowed.

7. What are the main productivity gains you get/envision your organization getting from use of generative AI?



SURVEY RESULTS

Business leaders expect greater productivity gains than cybersecurity professionals, but both say automated repetitive tasks are their lead usage, at 67% and 58%, respectively, and 62% for all respondents. Speed of production/service/results analysis is second, at 65% and 52%, and 59% overall. Routine and administrative tasks are third, at 58% and 45%, respectively, and 52% overall.

Forty-one percent of security professional choose “Write policies/courses, e.g. for security awareness/ training/education” compared to 38% of business leaders. And one cybersecurity professional correctly commented, “prefer you say ‘draft policies’ vs. ‘write policies’ and in general switch to the concept that it’s assistive but a human is still responsible.”

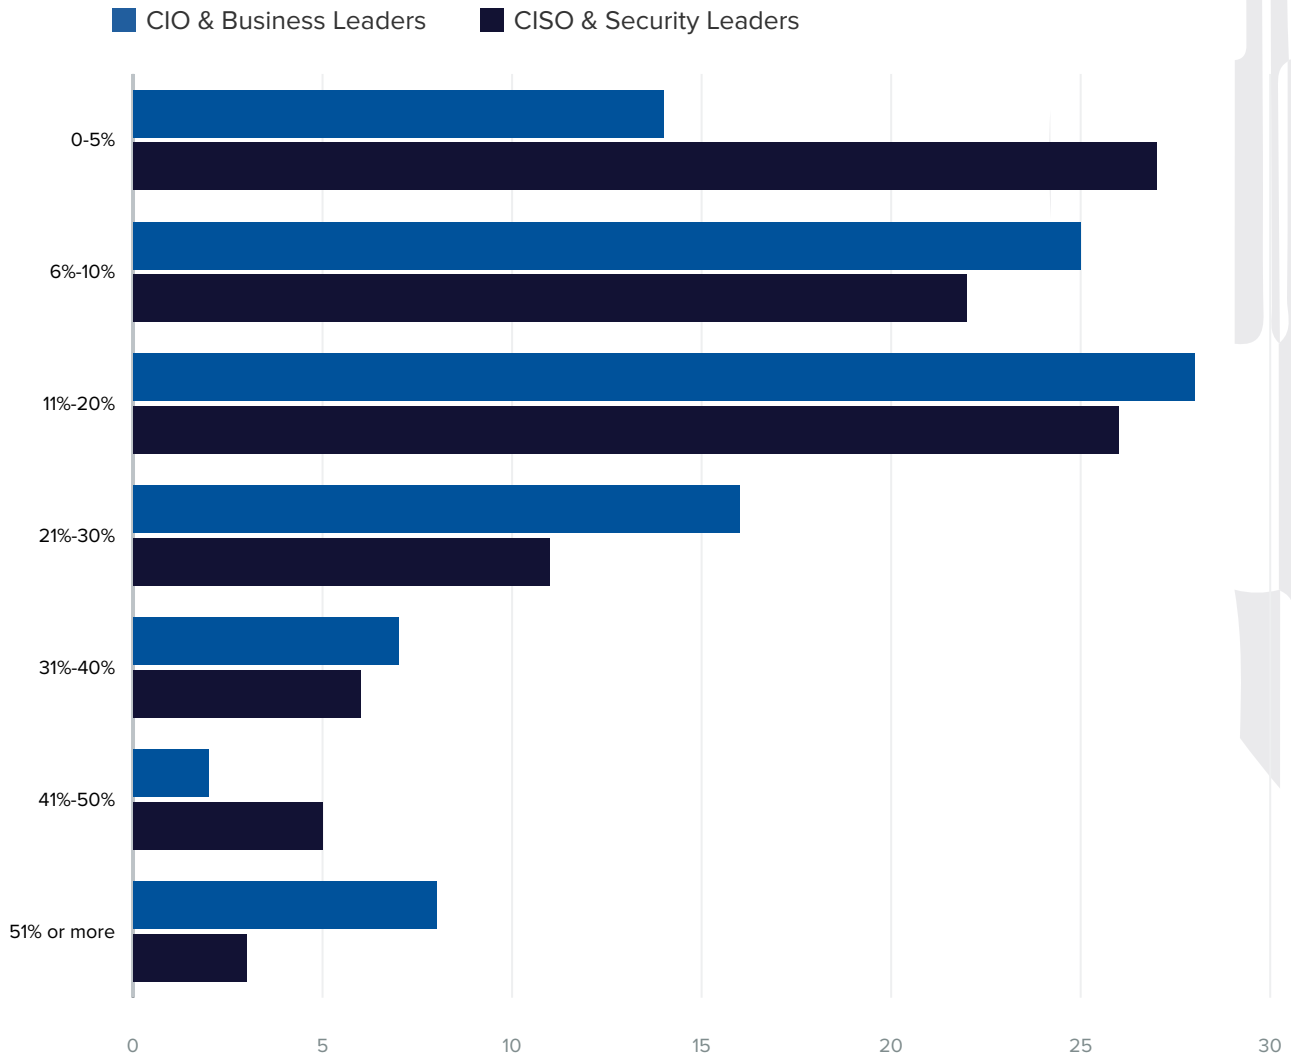
It may seem surprising, but reducing staffing requirement is not mentioned until sixth, and then only by 24% of respondents.

Reducing non-staff costs/budget, at 18% total, is more of a concern for business leaders, at 23%,

compared to cybersecurity professionals at 13%. Conversely, strengthening our own defenses, including choosing better passwords is more of a concern for cybersecurity professionals at 20%, compared to just 10% of business leaders.

Comments indicate that there are AI skeptics in both groups. One business leader says, “I don’t trust generative AI to produce anything without human supervision yet,” and a cybersecurity professional describes generative AI as more of a risk than a benefit.

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



The most frequently reported productivity gain is 11% to 20%, cited by 27% of respondents. More than twice as many business leaders - 8% - than cybersecurity leaders - 3% - reported productivity gains of more than 51%.

Among the 20% of respondents who report gains of 5% or less, cybersecurity professionals are over-represented at 27% compared to 14% of business leaders. This could be due to the

increased likelihood of non-deployment of generative AI due to operational restrictions, plus fewer use cases, since administration, sales and marketing – significant leaders in AI deployment – more often fall into the business leader category.

Where AI is implemented, productivity gains are significant, but business leaders report higher gains across a wider range of tasks.

9. For what use cases/ environments do you use/ envision your organization using generative AI?

There is a complete divergence in the answers from both groups, which is entirely understandable given the wide-ranging remit of business leaders compared to the more focused remit of cybersecurity leaders.


Among business leaders, the leading envisioned future use cases are preventing fraud at 84%. Legal/regulatory compliance is ranked second at 80% and medical diagnosis and treatment and medical results analysis are tied for third place at 75%.

Among cybersecurity professionals, the top envisioned future use cases are a tie in first place for medical results analysis, e.g., imaging and medical/pharmaceutical research, both at 90%. Second is medical diagnosis and treatment at 85%, and third is legal/regulatory compliance at 79%.

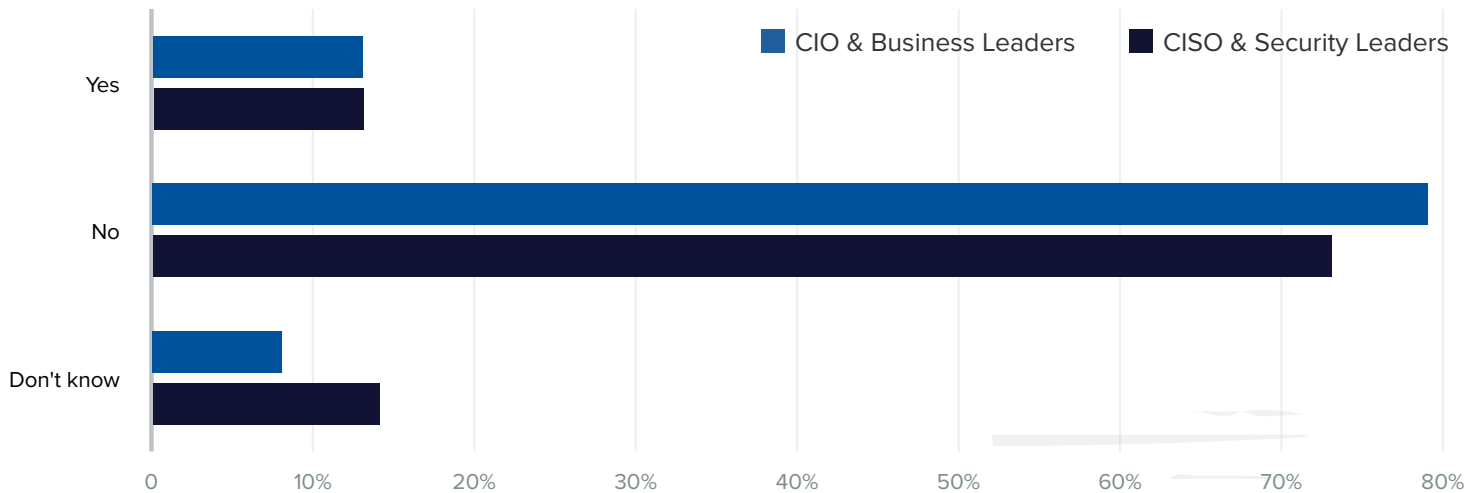
FULL CHART ON THE NEXT PAGE

SURVEY RESULTS

CHART 9 ■ CIO & Business Leaders ■ CISO & Security Leaders




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



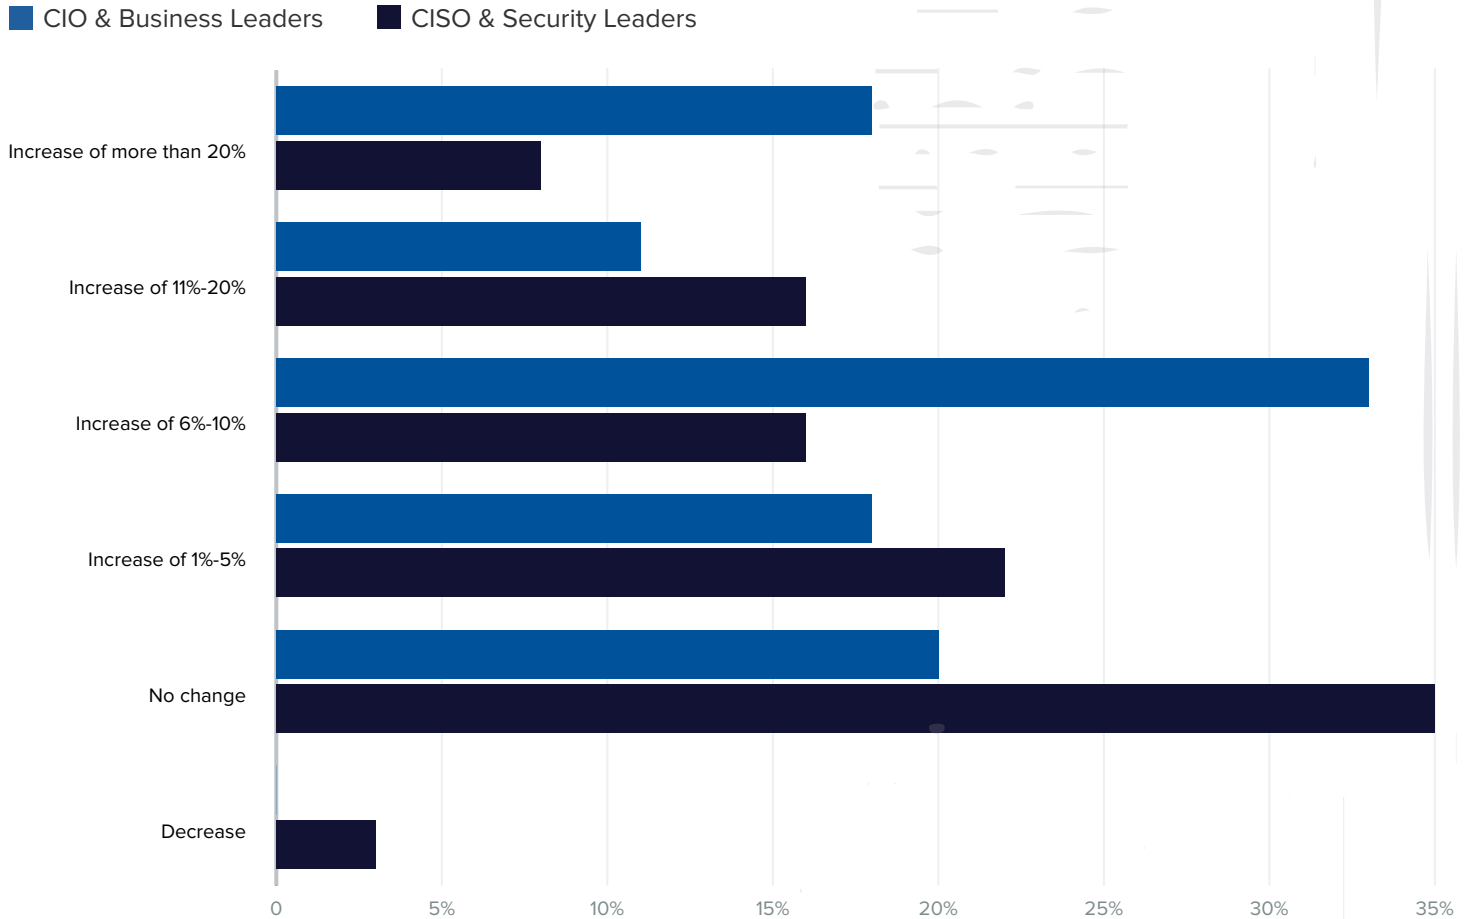
Thirteen percent of both groups have a specific budget for AI solutions, but more cybersecurity professionals - 14% - don't know, compared to business leaders at 8%. It is likely that generative AI's arrival is too recent, and there were too few commercially available use cases in April 2023 for generative AI to have made it onto annual budgets in the 2023/24 financial year. In addition, with no agreed allocation of authority over AI projects, they may not appear as a stand-alone item but could potentially be included in existing services that increasingly incorporate generative AI.

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




Sixty percent of business leaders and 49% of cybersecurity professionals expect to have a dedicated AI budget over the next year, which would represent a quadrupling of organizations with specific AI expenditure.

12. If “yes,” what % increase in budget for generative AI solutions do you expect in 12 months’ time?



Twenty percent of business leaders and 35% of cybersecurity say there will be no change in their budget. None of the business leaders foresee a reduction in their budget for AI, but 3% of cybersecurity professionals do.

13. Do you have specific plans to purchase AI-driven solutions over the next 12 months for any of the use case options earlier mentioned?



Thirty-eight percent of business leaders and 24% of cybersecurity professionals say yes. The number who don't know is also high, at 20% for business leaders and 38% for cybersecurity leaders.

14. If “yes,” please list up to top 5 desired use cases generative AI will address.


Responses from business leaders include security detection and prevention, marketing content creation, marketing automation, sales decision support, sentiment and behavioral analysis, back office productivity, media post-production - speech to text, tagging, and image generation.

Responses from cybersecurity professionals include asset management and patching, vulnerability management, legal and regulatory compliance, SOC operations, effective business continuity management, risk management, incident management, coding, marketing and other communications, report writing, research, diagnosis and treatment of medical conditions, speed for code writing, newsletters and blog publishing.

Cybersecurity professionals also list use cases for chatbots for customer support, language translation and localization, more accurate and context-aware language translation, art and design, and software development.

Although cybersecurity leaders were less likely to have specific purchase plans than business leaders, the 24% who did - see Chart 13 - had a wider range of specific planned purchases than business leaders.

15. What are your main concerns when it comes to implementing generative AI by yourself and/or by others?



Although there are differences between the two groups regarding concerns about particular threats, the top concern for both groups is leakage of sensitive data by staff using AI, cited by 80% of business leaders and 82% of cybersecurity professionals.

Second for both groups is ingress of inaccurate data - hallucinations, cited by 71% of business leaders and 67% of cybersecurity professionals.

In third place for both groups is AI bias/ethical concerns cited by 61% of business leaders and 57% of cybersecurity professionals.

16. What do you view as the biggest risk within code repositories in cybersecurity when it comes to generative AI use?

Business leaders are most concerned about loss of code. They also mention unintended consequences and privacy concerns; the embedding of malicious or dysfunctional code; misuse by bad actors, e.g., deepfakes or misleading information; loss of confidentiality; ethics/bias; ransomware; and phishing.

Cybersecurity professionals are most concerned about visibility of where code comes from, i.e., Is it proprietary, open source, poisoned or malicious? They also share many of the concerns of business leaders, including introduction of malware or copyrighted source code, skills loss, information getting into the wrong hands, and information code getting corrupted.

Comments on this question include: The employee need enough skills in the area to know when the AI is hallucinating or returning bad code. There is also concern about staff not understanding the code but using it because it works, leaks of data and ransomware created by generative AI, deepfakes, copyright issues, access; leakage of sensitive and proprietary data, difficulty in auditing the actions of individuals vs AI, and incorrectly configuring/ implementing AI products.

Then comes the deliberate misuse of AI products, poisoned concepts or poisoned inference of decision vectors; algorithms that don't actually work but appear to and accidental usage of open-source code in proprietary code creation.

One respondent says certificate management, which is already very hard to do well, will become critically essential to maintain confidence. Most organizations

are not ready to do certificate management even poorly let alone at the level required to provide assurance in data and systems.

Another comment: "We're starting from the business end with AI and haven't yet considered generative AI's access to code repositories ... that'll come late next year at the earliest."

17. What tools, processes or approaches do you currently use and intend to use to mitigate the concerns around use of AI by your own organization or your supply chain or partners?

Seventy-three percent of business leaders and 69% of cybersecurity professionals currently use AI for encryption of data.

Fifty eight percent of business leaders and 48% of cybersecurity professionals currently use AI for psuedoanonymization of data.


It is significant that 38% of business leaders and 48% of cybersecurity leaders intend to continue banning the use of generative AI in the workplace and that 73% of business leaders and 78% of cybersecurity professionals intend to take a walled garden/own AI approach going forward. Both suggest a return to the wall and moat of the past as businesses strive to regain control of the AI genie that has been let loose from its bottle.

In comments, one business leader says: “We have a policy on the use of generative AI in place,” and one cybersecurity leader says: “Currently - no controls in place or planned until after something bad happens to peers.” Another says, “While currently banned, gen AI will be governed by policy requiring human intervention/review of any generated work product.”


FULL CHART ON THE NEXT PAGE

SURVEY RESULTS

CHART 17 ■ CIO & Business Leaders ■ CISO & Security Leaders

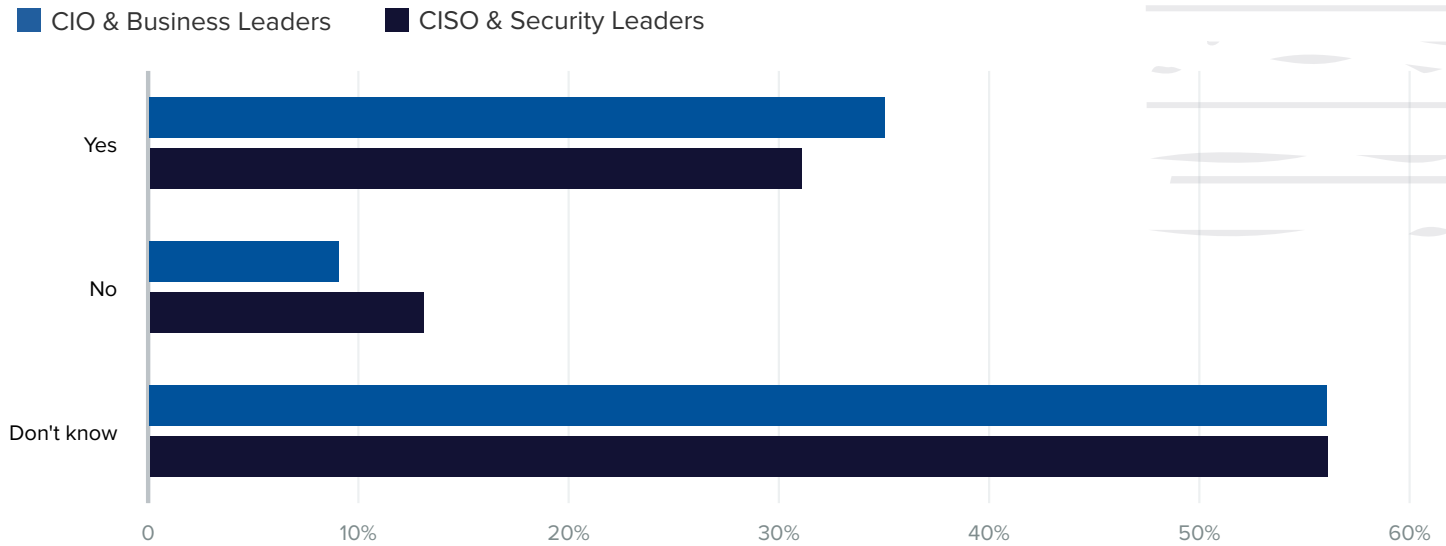


18. Is there a process/ playbook/guidelines/ policy in place to ensure that all generative AI usage/deployment in your organization complies with agreed security policies?



Thirty percent of business leaders and 31% of cybersecurity professionals say that they do have playbooks for AI deployment.


19. Do your competitors currently use generative AI?



Thirty-five percent of business leaders and 31% of cybersecurity leaders say their competitors use generative AI. An exceptionally large number of respondents - 56% - say they do not know.

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

■ CIO & Business Leaders ■ CISO & Security Leaders



A worryingly low 38% of business leaders say they do understand these regulations, as do 52% of cybersecurity leaders. Yet, given the pace of change and the lack of global standard regulations, this is perhaps not surprising.

Standout Survey Results

TONY MORBIN: What particularly stood out for you in the results, and what's your take on that?

ANTON CHUVAKIN: I saw some adoption anomalies as I was reading the report, but the report made sense in most cases. The contradictions between security and leaders made sense, but some of the adoption numbers or perceived adoption numbers looked really high. These are maybe slightly biased.

STEVE POVOLNY: The results of the report are pretty on point with what we see in industry. Some of the largest standout surprises were the discrepancies between business leaders and cybersecurity professionals.

DAVID BAILEY: I was pleased to see the different respondents from the types of leaders within the organization. It's good to want business leaders to be able to utilize technology to be successful, and AI is going to help that, which is great. The downside of that is the concern and apprehension from security professionals as well as those that need to manage risk within the organization. But I'm glad that some of that apprehension is there because there are a lot of unknowns yet to be decided on how organizations have to manage their risk.

LAURENCE MCNALLY: The survey results correlate to what I'm seeing as I talk to businesses on using AI. Business leaders are more bullish as opposed to our cybersecurity folks who are definitely more skeptical and thinking about the trustworthiness and side effects of the AI. Another thing that stood out to me was the number of people that said they understood the regulations.

Why Banning AI Usage Won't Work

MORBIN: Quite a few respondents, particularly the cybersecurity professionals, say that they were banned from using AI in their organization. Is banning the use of general AI for employees or the business an effective way to mitigate threats?



Anton Chuvakin

Security Adviser at Office of the CISO
Google Cloud

Chuvakin is an expert in log management, SIEM and PCI DSS compliance and is an author of the books "Security Warrior," "Logging and Log Management" and "PCI Compliance, Third Edition." He was previously research vice president and distinguished analyst at Gartner for the Technical Professionals Security and Risk Management Strategies team. He was also a director of PCI compliance solutions at Qualys and worked at LogLogic as a chief logging evangelist.

There are two types of AI systems. There is the LLM world that is not really high-risk, but these other applications, like regression models and all of these other vision models, are very different from the LLM world.

- Laurence McNally

POVOLNY: We know how things work when you ban holistically or make a broad strokes approach like a ban: People find ways to work around it. This is one of the most commonplace and polarizing issues around generative AI – how to use it appropriately. Do we get aggressive with something like a ban? This is a personal decision and a business decision, and it's hard to be too judgmental of either of those. It can be an effective way to mitigate risk holistically.

But on the flip side, employees will actively find ways to work around it, which can be more damaging than just training them effectively how to use it or limiting, controlling or having some oversight on the approach to usage. The FUD – the fear, uncertainty and doubt – surrounding generative AI shouldn't be a reason to holistically ban it. We should control and educate and enforce the usage of it effectively.

CHUVAKIN: Bans ultimately cause usage to increase – sometimes in all sorts of insecure ways. I'm against banning because ultimately, banning often produces the opposite effect.

MCNALLY: At the companies that I was working with that banned ChatGPT, other tools such as Aha were using it, so people were using that. A ban just pushes the problem down to somewhere else.

Guidelines for AI Usage

MORBIN: Is there a lack of guardrails for the use of AI because people don't know what the best options are or because they don't have the skills to implement them? Or is the issue of security just not high enough up the priority list compared to getting the benefits of being an early adopter?

BAILEY: One of the foundations of a really strong security program is to ensure that you've got good governance, guidelines and standards. Security is not just an IT problem or a security problem; it's a business problem. While organizations may have business leaders that want to embrace the use of AI, they do not yet have in place the right governance, the right stakeholders identified or the right understanding of what it takes to address the impacts of AI – the trustworthiness and the risks associated with it – and then implement that throughout an entire system or software development life cycle. A lot of the organizations we deal with are struggling to just get the maturity that is required for today, let alone using AI.

The guidelines for organizations ultimately will come down to: Do you have the mechanisms in place to know what the risks of using AI are, and do you have the people and processes in place to address it? Some data scientists are excited about using AI for outcomes and look at AI as an enabler of their process, but some security professionals

look at it as a disruptor. They know AI is not going anywhere and that they are going to have to embrace it, but they are concerned about all of the things that are required to do it in a way that is reasonable and appropriate from a security standpoint and a risk standpoint. We're going to have to develop some new processes in order to make sure we're doing that effectively.

Securing AI

MORBIN: What exactly can we do to mitigate the risk of generative AI being used for malicious purposes?

CHUVAKIN: This is my main goal. We recently published a paper called "Securing AI: Similar or Different?" which answers some of these questions. Let me give you a broad framework. First, some of my colleagues rush to thinking that to secure AI you need AI. In reality, one of the guardrails may be improved data governance. Some of the recent breaches involving AI, including losses of training data, had nothing to do with actual AI workloads; they had to do with processes related to training data being broken.

Think about whether the controls that you have always had and used are relevant in their intact form. For infrastructure security, if you are securing where you prepare the data or where you run the AI workloads, this applies verbatim. ChatGPT or Bard or commercial enterprise-type AI solutions are ultimately software-as-a-service products, so much of the SaaS security applies. This bucket is called "Ultimately, there's no difference." Some of the controls are the same.

But there's also a more exciting bucket called "These controls are different," and these controls maybe have different emphasis. For example, think about data filtering. If you have a massive CRM application, a traditional data-intensive

enterprise app, you filter data. You want to not have malicious data coming in, but ultimately whatever comes out is only what you put in. With AI, what comes out is not what you put in. It may be something else. So, filtering inputs is a great idea. But filtering outputs is new. That's an example of how data security control morphs quite a bit when you add AI.

If I think of data governance, I think of decisions to be made more tightly coupled to the data life cycle, like, "What data goes into training? How do you secure prompts? Who can see the prompts?" All this is a blend of traditional and novel controls. With threat detection and response, if you stick to security scope, there are some changes, but not dramatic changes. But when you start thinking about the content safety, a whole world opens up that you may not have encountered as a CISO. You've dealt with threats, badness, hackers and insiders, but you haven't dealt with machine-produced content that harmed your company. Now, you need to think about it. The CISO team's responsibilities expand to areas that they're not familiar with.

The final example I'll give on the controls is: Some people say that the number one problem they have with AI is intellectual property. My reaction is, "But your job is security, right? You are a CISO. Why is this your problem? Can you shove it to somebody else's inbox? And the person says, "Guess how I ended up with the problem? Everybody shoved it off their inboxes, and it ended up in my inbox because it vaguely connects to risk." We have to solve these problems, but they're unfamiliar problems. Traditional security teams don't know what to do about [securing AI]. That is an exciting challenge. The expansion of the mandate is what freaks a lot of people out - not that they have to deal with adversarial prompts.

BAILEY: I totally agree. It's important to focus on data governance – understanding what data you have and then knowing how AI will impact that data. Most people put data in and then want to interact with it to get an outcome. Well, the outcome may be completely new, and that requires determining trustworthiness, potential harm and potential impact. We may have to adapt new things for existing processes in order to affect that good data outcome. Data governance is extremely important in your AI journey.

POVOLNY: A lot of the threats that we think of surrounding AI in general as a concept aren't fundamentally new in the way that we protect and monitor data. Data protection extends to protecting your models and your training data from poisoning. Data validation and explainability are very similar to code reviews and code auditing. A lot of techniques that we know already just have a new application here.

Cybercriminals are going to find ways to deploy and exploit AI-based attacks regardless of how well we do that, so when we can simulate research and have a deep understanding of what those attack methods look like, it really helps us to identify and determine what the tools and techniques will look like when we see them in the wild. This is one of those rare times as an industry that we're on equal footing with the cybercriminals. We're just as far into the research and development of techniques and applications as they are for the malicious counterpoints. We have at least an even footing there, if not a step up, and that's exciting.

MCNALLY: Even outside of cybersecurity or cybercriminals, when your own data scientists are putting data into the model, especially LLMs, people are exposing all their Confluence documents without looking through and doing data discovery and redaction. Then, they're



David Bailey
Vice President of
Consulting Services
Clearwater

Bailey works alongside men and women for the only company combining deep healthcare security and compliance expertise with comprehensive service and technology solutions to help organizations become more secure, compliant and resilient. He serves integrated delivery networks, digital health companies and the defense industrial base in achieving their missions.

surprised that sensitive information is coming out of the model. There should be a gatekeeper between what gets fed into these LLMs and whether you put any security keys or tokens into the model. If you put all Confluence in and there were some security tokens in that, the LLM model can give an output of the security token.

POVOLNY: The risk of poisoning your own models is higher, or at least equal, internally as it is externally.

Regulations for AI

MORBIN: Fewer than 40% of business leaders in the report say they understand the regulations relevant to their geography or industry. You may be skeptical of even the 40% figure, but how can organizations catch up, and how can we ever hope to have globally agreed regulations, given cultural perspectives on privacy and security and where the balance is? *Note: This discussion took place prior to the U.S. Biden executive order on AI.

MCNALLY: It's a really hard question in terms of the agreeability of all the different regulatory bodies. I'll keep that piece out because that's a very long rabbit hole. But in terms of a business leader trying to adhere with whatever regulation that they choose that they want to adhere to, I go back to the example of GDPR. It was one of the trendsetters for the privacy space, and we see this happening again with the EU AI Act in Europe. So, getting companies up to speed and getting prepared for the EU AI Act is one piece where they can get ahead of the curve.

In terms of risk level, not all AI is of the same risk. If you have an employee using ChatGPT or Bard internally to help them draft an email, that's very different than an AI system that's predicting loans or being used in healthcare. That's way more risk. We can help organizations build out an inventory, rank the riskiest AI systems, go after the high risks and put the regulations or policies and procedures on those high risks.

CHUVAKIN: I am super skeptical about the respondents saying they have full understanding of regulations because I don't think that's the case. We have a team that tracks regulations affecting AI, and they're about to overflow the spreadsheet maximum row number with all the entries. A lot of stuff is being reapplied or refocused on AI, and the future is going to be very freaky.



Laurence McNally

AI Governance, Data Discovery and
Technical Product Manager
OneTrust

McNally leads governance of AI products and is responsible launching the latest at OneTrust, specifically for OT Global Data Platform Products that consists of internal products (developer microservices) and consumer-facing products (Low Code No Code LCNC Platform).

I don't know if small startups will build it for their own regions and then hope for the best. I don't know how we're going to deal with that, especially when it comes to contradicting regulations. Which one do we follow? The lack of understanding combined with high speed of adoption is a hugely explosive combination. I have no idea what will happen in this area, and I don't know anybody who does.

MCNALLY: Regarding the right for your personal data to be forgotten, once a model has been trained, deployed and shipped, your data is in there. The right to be forgotten goes away. You can't put in a DSA request and expect them to train the model again. That is a whole other rabbit hole of technicalities there.

POVOLNY: To be fair, we lost that right with the advent of social media as well. Privacy is a complete fallacy nowadays. But this is another application of it.

Top-Priority Concerns

MORBIN: Are the priorities that the respondents express around their concerns broadly in line with what you or your organization sees as the most important risks, or what do you see as the most important risks?

POVOLNY: The concerns about employee data and company data making it into models, about the way that attacks are being deployed and used, the strengthening of common and legacy types of attacks, such as social engineering and phishing, are obviously being dramatically improved through some of these tools. All of those hold true and are some of the risks that we see inherently.

This is an explosion of technology much in the same way as the development of the iPhone, or maybe the personal computer. There is going to be a red-hot period where the world innovates and decides how they're going to use and explore and push the boundaries of AI. Even though it's 70 years old as a concept, it has a rebirth now. It's no-holds-barred, but we have to be prescient about what the applications and risks are. We have to think about how to control them and apply them without putting handcuffs on the capabilities of AI.

MCNALLY: An interesting point from the survey was the consensus around the leakage of IP. In two questions, people say leaking IP of a company is one of their main concerns. A company's data uses so many different vendors, you don't know what data of yours they're using to retrain. Jira and Aha have introduced generative AI within their applications. Are they using our documents to train their model that's being shared with an organization? Even worse: Are they using some of our customer data? In the example of Salesforce Einstein, is our CRM being used to train Einstein, which is being shared with other organizations? That threat goes beyond an employee going on ChatGPT and putting in something that they shouldn't. When it involves the vendors that you're using, there's a huge level of risk there.

CHUVAKIN: In the survey, sensitive data leakage is number one, ingress of inaccurate data hallucinations is number two, and then the third bucket is broad bias/ethical concerns. And that make sense. The only slight change is the ingress of copyrighted IP. For some reason, they're talking about ingress, not egress. IP in certain copyright being produced is not coming up in the surveys. Google just announced indemnification for the enterprise AI models. It comes up a lot, and it doesn't come up at all in the survey. It's not about your IP showing up in the AI; it's about whose IP is the stuff that the AI produced. If somebody points at it and says, "Hey, I recognize this code. I wrote it," then suddenly problems happen.

Generative AI and Healthcare

MORBIN: David, looking at organizations that Clearwater works with within the healthcare sector, what are the concerns there? Is generative AI already being deployed in medical applications, and how do they manage to do that given the potential liabilities in that sector?

BAILEY: There is a level of awareness at the industry level. The industry understands and knows that AI is here. We're dealing with many organizations that are struggling with the knowledge that they have to implement the governance aspect. Healthcare today is all about the patient engagement, patient experience and clinical outcomes. AI applies well to patient engagement, patient experience and productivity, and you can see where the vendors can utilize productivity and outcome. When you're dealing with true medical application, you're at the bedside with the patient and you're at some level of use of AI for clinical outcome, there's still a lot of concern about trustworthiness and knowing how to address the right outcomes.

In research, AI is being used in imaging to address and look at images, process images, find tumors and scan. There's so much applicable use. The full-level adoption is not there yet, but the concern is real. Organizations will struggle over the next year or two to ensure that they have the right stakeholders and processes in place and that they can look at what that outcome is, especially from a clinical outcome perspective, and know that they can trust the outcome to make good decisions for their patients and use that technology with good clinical care in mind.

MORBIN: It's the difference between strategy and operations. Our cybersecurity professions, the people who have to implement it, have a bigger struggle than our business leader



Steve Povolny
Director of Security Research
Exabeam

Povolny has more than 15 years of experience leading global teams of security researchers, data scientists and developers. At Exabeam, he and his team have a singular focus: integrating world-class research into the industry's top cybersecurity solutions to disrupt cybercrime and defend customers' critical assets. As a regular speaker at industry conferences, Povolny shares insights on emerging trends, attack surfaces and cutting-edge vulnerability and malware research.

respondents, who are talking largely of intended use or expected use. A lot of them put the medical applications very high up on their list, but implementation is a little bit harder.

BAILEY: We're seven to 10 years into a network-connected medical device, and it has been a struggle to ensure that there is an appropriate level of security and reasonable and appropriate controls with network-connected medical devices, knowing the threats that exist on the network. So now, when you add generative AI, learning models and machine learning to the process, we've got a long way to go. There's a lot of risk to identify and mitigate.

MCNALLY: The EU AI Act is aiming to bucket these systems into four different categories, starting with unacceptable. Unacceptable things include something that's a threat to someone's safety or livelihood or the rights of people. Companies can't build that. Under the EU AI Act, all of the healthcare uses would be seen as high-risk. That's why the inventory of AI systems is so important. The company needs to know all the systems that they have in place and the models that roll up into that.

The EU AI Act is aiming to bucket these systems into four different categories, starting with unacceptable. Unacceptable things include something that's a threat to someone's safety or livelihood or the rights of people. Companies can't build that. Under the EU AI Act, all of the healthcare uses would be seen as high-risk. That's why the inventory of AI systems is so important. The company needs to know all the systems that they have in place and the models that roll up into that.

AI for Vulnerability Discovery, Mitigation

MORBIN: One of the biggest issues with AI is trust, with hallucinations potentially impacting the validity of results. Twenty-three percent of respondents say they are using generative AI to find and fix vulnerabilities, and Steve has said he is skeptical about implementing vulnerability discovery and mitigation via generative AI. Steve, please explain that.

POVOLNY: I'm skeptical of the concept of 23% of respondents truly discovering and mitigating code-based vulnerabilities in any kind of automated and effective fashion using generative AI. What I'm not skeptical of is that there is probably a frequent use of code review and basic bug fixes and development processes that generative AI can aid in where classical software bugs and configuration issues are likely possible to be discovered and mitigated. We're seeing research leading the effort, but this is very much prior to any kind of market application in things like zero-day discovery, deep reverse engineering of code, and complex bugs that still require a lot of human intervention and human knowledge to discover. So it's probably more about a broader set of terminology around vulnerability, discovery, and mitigation.

The Need for Human Intervention

MORBIN: Others, what can we trust or where do we need to get humans involved?

CHUVAKIN: For a lot of answers we want, AI will give you a candidate answer, but if you treat it as the right answer every time, you're going to go very badly and spectacularly wrong. Ultimately, human skills are very much needed. But if you use these AI models for ideas – for things to try, things to do, candidate answers – they're really good.

MCNALLY: There are two types of AI systems. One is the really cool applications that have democratized AI to users to help them draft an email and help them with ideas. That's one sense of AI, and there are regulations with that. But the systems you're talking about are really complicated, with core data scientists involved and very different procedures and policies. There is the LLM world that is not really high-risk, but these other applications, like regression models and all of these other vision models, are very different from the LLM world.

POVOLNY: That's a super important distinction to make: Is there a fundamental difference between generative AI, which is the creation of computer-driven or computer-aided content in some form of media, at least in most uses are today, and traditional AI and ML, which might be GANs or AGNs or the creation or recognition of content, pattern recognition and creation, and classification algorithms. These things don't tend to overlap, but they do get conflated in the concept of generative AI versus AI in general. We need to be super careful when we use these definitions that we don't overlap them.

Other Survey Results

MORBIN: Did any of the other results we haven't mentioned so far stand out or surprise you?

BAILEY: For the one in which 31% say that they already had plans to purchase AI-driven solutions in the next 12 months, what is an AI-driven solution? You would hope that 31% had gone through some level of risk analysis and understanding of what that means, what the risks and impacts are to the organization, and how it feeds into the entire business impact to that organization. It's great that you can go buy some AI-driven system, but how it fits into the whole life cycle and trustworthiness and risk acceptance is where we're lacking.



CHUVAKIN: The real surprise is not just the high usage, but the question: For which use cases do you either use or predict the use of AI? Legal and compliance is at 80% with current use at 20%. So while I can understand the desire to use, say, an LLM-based summarizer to understand certain inscrutable compliance mandates, I can imagine a very safe, very auditor-proof, very tame usage for compliance use cases. I have the deep suspicion that's not what they mean. I have a suspicion that they're going to answer compliance questionnaires with LLM bots. They're going to write destination statements with machines and with light review of humans, and that's going to produce incredibly fun-to-watch disasters for them.

When lawyers tried to argue cases using ChatGPT reasoning, it ended up being 90% faulty and based on made-up data. So, the compliance usage predicted at 80% of all respondents is exciting, fun and probably very failure-prone, and to me it is a surprise.

POVOLNY: One of the things that really stood out to me was that most of the survey respondents indicate that the businesses think the C-level staff is responsible for deploying and maintaining generative AI solutions. That makes no sense to me, except when you think about it, the C-level staff is ultimately writing the check and is responsible for the strategy behind it. We're going to see an evolution as companies start to realize that they're missing skill sets and capabilities in the data science realm. They'll need to make sure that they have a chief data science and a data science organization that can effectively deploy and maintain these solutions, obviously rolling up to the C-level staff.

MCNALLY: At the very start, I mentioned the overall bullishness of the business leaders to use these solutions versus the concerns of cybersecurity. That stood out to me because the use of AI across businesses is so distributed. You have different teams using different ML ops tooling, and you have employees using vendors that might buy some shadow IT that has AI. You have a distorted view across the whole landscape of an organization. If you go into an organization and ask, "Where are you using AI systems and why?" they can't quickly pull out a report. There's a lot of confusion among the C-level executives and the higher-ups on where AI is actually being used and why it's being used.

Securing AI for business is a long slog. There is no magic in this area; you just need to work hard and learn it and then secure it.

- Anton Chuvakin

The Future of AI

MORBIN: What are your predictions for the future of AI and security, particularly whether it's going to be more of an ally for security or a threat?

BAILEY: I'm not a "sky is falling" security professional. I try to apply reason to this. Organizations that are not in front of this train will get run over by the train, and it's important for organizations to focus on this now. AI is here to stay, and we have to start addressing it.

MCNALLY: The different types of AI, your regression models and legacy AI, won't change in the next couple of years. The hallucination space of LLMs, that fearmongering, will reduce. You Google something today and if it's not the right article, you use your common sense to figure out what's right and not; you don't just take everything verbatim. The risk of models hallucinating will die down a little bit. A lot of the scare around deepfake images is warranted – what's actually AI-generated content versus what's not? So, we should have labels or some system that has to add them. But then, there are ways around that too. But I don't buy into the idea that the sky is falling because of it. It's a net positive on productivity.

POVOLNY: The misinformation is the biggest risk that I see coming out. We'll have to have systems in place to identify, defense in depth, validation, and additional checks to ensure that the content that we're consuming is actually the content that we think we're consuming. The world is badly trained on that front, and generative AI is going to make that problem more difficult – no question about it.

But I'm definitely on the ally side of things. I think it's going to be revolutionary, already is revolutionary. The applications have a profound impact on nearly every industry vertical worldwide,

and the pros will outweigh the cons so long as we can get past the idea that it's a silver bullet that fixes everything and find out where the real applications are.

CHUVAKIN: Our CISO, Phil Venables, makes a good argument that ultimately, in the long run, AI will favor defenders not attackers because ultimately defenders are the side with more data. That generates a lot of very exciting optimism for using AI for security because if the technology revolution inherently favors defenders, the security of other things will improve because of AI. It is a useful prediction to say that AI favors the defenders over attackers because of the amount of data, but what about the other side – securing the AI used for business for other purposes?

That prediction is a long slog. Let's say we're going to secure mobile. Likely in 10 years, we more or less know what we are doing. Despite all this noise about AI, we do see companies that just encountered cloud for the first time and they're marveling at how different things are in the cloud. For them, the revolution of securing a new venue is now, but the venue is cloud, not AI. And we roughly know what will happen. They will go through a journey and normalize their relationship with this new terrain to secure.

We are in the beginning of that journey for AI. We know what to do. We know which data security controls are more relevant. We know what governance tricks work. We know how to detect and respond to new threats. But ultimately, securing AI for business is a long slog. Some people, like Google, will be there first, but many others will encounter AI for the first time in 10 years. That's my prediction. There is no magic in this area; you just need to work hard and learn it and then secure it.

CONCLUSIONS

The use of generative AI is expanding, and so are expenditures for it.

Utilization of generative AI is exploding, and though only 15% of respondents are actively deploying AI, when those conducting trials or planning to implement it are included, the figure reaches 70%, hence the high growth projections.

Expenditure specifically on generative AI is multiplying rapidly. Our research shows that the numbers reporting specific budgets for gen AI is set to increase fourfold, and budgets allocated are expected to increase by 10%, however it is likely that these are minimum figures.

Surprisingly, only 38% of business leaders, and even fewer cybersecurity leaders - 24% - have specific plans to purchase AI for any of the use cases covered. The difference reflects both the more cautious approach of security professionals and the wider range of deployments expected by business leaders. Consequently, a significant proportion of businesses expects multiples of growth in expenditure in a technology where they are not sure what they will buy or how they will use it when they do. But they expect to buy it so anyway.

The growth in deployment and expenditure is expected to be much higher than even our respondents' projections as the introduction of new gen AI use cases and their increasing familiarity and proven productivity gains both see wider and deeper adoption. These expenditure growth figures will be further masked by the adoption of generative AI within the tools and services of existing suppliers.

Use cases for generative AI are growing, and so is productivity.

What is clear is that, notwithstanding concerns around security, privacy and safety, generative AI represents a paradigm shift in how business works, and it is currently seeing unprecedented accelerating adoption. This is being driven by our business leaders who are experimenting across a wide range of gen AI tools and a plethora of use cases. While cybersecurity leaders are more cautious, they too recognise the gains and are experimenting, albeit in a narrower range of use cases and tools.

The productivity gains exceed 10% in most cases, though they appear to be higher for business leaders than cybersecurity professionals.

Both business leaders and cybersecurity professionals are aware of the potential pitfalls and are largely in agreement about the prioritization of potential negative consequences of inherent flaws, accidental or deliberate misuse. In particular, data loss, ethical concerns/bias and ingress of inappropriate/poisoned data need to be prevented/mitigated, and cybersecurity professionals tasked with achieving this tend to prioritize the need for security above the need to improve productivity.

Generative AI is still being banned, and a walled garden approach is coming.

Approaches to mitigating threats vary, and outright bans on the use of generative AI are more common among cybersecurity leaders, but the number advocating an outright ban on use in their organization is surprisingly high. Thirty-eight percent of business leaders and 48% of cybersecurity leaders expect to continue banning the use of generative AI in the workplace – which contradicts the 70% planning to use AI. Also, this approach is not considered viable by many in the industry, as our expert analysis shows, since it could replicate the “shadow IT” issue in AI as users circumvent the rules with less known and potentially less secure AI variants.

The need to address risks is also reflected in the statistic that 73% of business leaders and 78% of cybersecurity professionals intend to take a walled garden/own AI approach going forward. These approaches may create issues about limiting the ability of generative AI to learn, but the respondents did not name this as a concern.

Understanding of AI regulation is low.

Understanding of regulations in any particular vertical or geography is low, as 38% of business leaders say they do understand these regulations, and 52% of cybersecurity leaders say the same. Our expert panel feels that even these low figures are probably higher than reality given how quickly regulations are developing and the fact that they are not standardized internationally and are potentially contradictory.

Guardrails are needed.

It broad terms, it appears that business leaders understand that generative AI represents an unprecedented opportunity for increased productivity, and cybersecurity professionals see

Organizations that are not in front of this train will get run over by the train, and it's important for organizations to focus on this now. AI is here to stay, and we have to start addressing it.

- David Bailey

the unprecedented risks posed by generative AI. But at the same time, business leaders know the risks and the need to engage their cybersecurity professionals to mitigate that risk. And cybersecurity professionals recognize the opportunities generative AI deployment affords their company and their own profession, and therefore, the need to embrace deployment.

While the perspectives of business professionals and cybersecurity professionals differ, it appears that they are cooperating to implement guardrails to ensure productive and secure deployment of generative AI. But knowledge and understanding of how best to do that has not been established when it comes to the details of what approaches will be most effective, and we are currently in a period of trial and error.

ABOUT THE SPONSORS

Google Cloud

Whether you're migrating or already in the cloud, we'll help you modernize and digitally transform your business. Build with generative AI, deploy apps fast, and analyze data in seconds—all with Google-grade security.

For more information, visit: cloud.google.com

Clearwater

Clearwater is the only company combining deep healthcare security and compliance expertise with comprehensive service and technology solutions to help organizations become more secure, compliant, and resilient. Propel your mission forward with the leader in healthcare cybersecurity & compliance.

For more information, visit: clearwatersecurity.com

Microsoft Security

Microsoft Security understands the evolving landscape of cybersecurity threats and the critical need for advanced solutions to safeguard your business. By harnessing the power of generative AI, Microsoft Security provides a proactive and adaptive defense mechanism that learns from vast datasets, anticipates potential risks, and evolves alongside emerging cybersecurity challenges.

For more information, visit: microsoft.com/security

exabeam®

Over 90% of breaches are rooted in compromised credentials and most security tools can't help. Exabeam offers a breakthrough combination of capabilities that security teams needs in products they will want to use. Advance your AI-driven security operations: detect threats, defend against cyberattacks, and defeat adversaries with Exabeam.

For more information, visit: exabeam.com

onetrust

OneTrust is the trust intelligence cloud platform organizations use to transform trust from an abstract concept into a measurable competitive advantage. Organizations globally use OneTrust to enable the responsible use of data while protecting the privacy rights of individuals, implement and report on their cyber security program, make their social impact goals a reality, and create a speak up culture of trust.

For more information, visit: onetrust.com

About ISMG

Information Security Media Group (ISMG) is the world's largest media organization devoted solely to information security and risk management. Each of our 28 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 and fraud. Our annual global Summit series connects senior security professionals with industry thought leaders to find actionable solutions for pressing cybersecurity challenges.

Contact

(800) 944-0401 • sales@ismg.io • research@ismg-corporate.io









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