Gartner
Peer Community™

Navigating Generative AI Governance: Infosec Leader Perspectives

Organizations are looking to generative AI (GenAI) governance as the tech's risks and opportunities continue to emerge. What are their strategies and how many plan to implement generative AI for security?

One-Minute Insights:

- Nearly two-thirds of surveyed leaders say their organization has or is developing a generative AI governance strategy
- Over one-third have no plans to offer training on proper usage of generative AI tools, but some have a corporate policy
- Nearly all respondent organizations are considering generative AI use cases for security, with chatbots and threat intelligence as common choices
- The most common generative AI vendor strategies among respondents include evaluating model transparency or privacy policies
- Most surveyed leaders plan to hire in the next 12 months to meet generative AI governance or security needs

One-Minute Insights on timely topics are available to Gartner Peer Community members. Sign up for access to over 100 more, and new insights each week.

Data collection: Aug 24 - 29, 2023

Respondents: 235 IT and information security leaders involved in security and/or risk management efforts related to the use of generative AI tools at their organization

One-fifth of respondent organizations have established a GenAI governance strategy

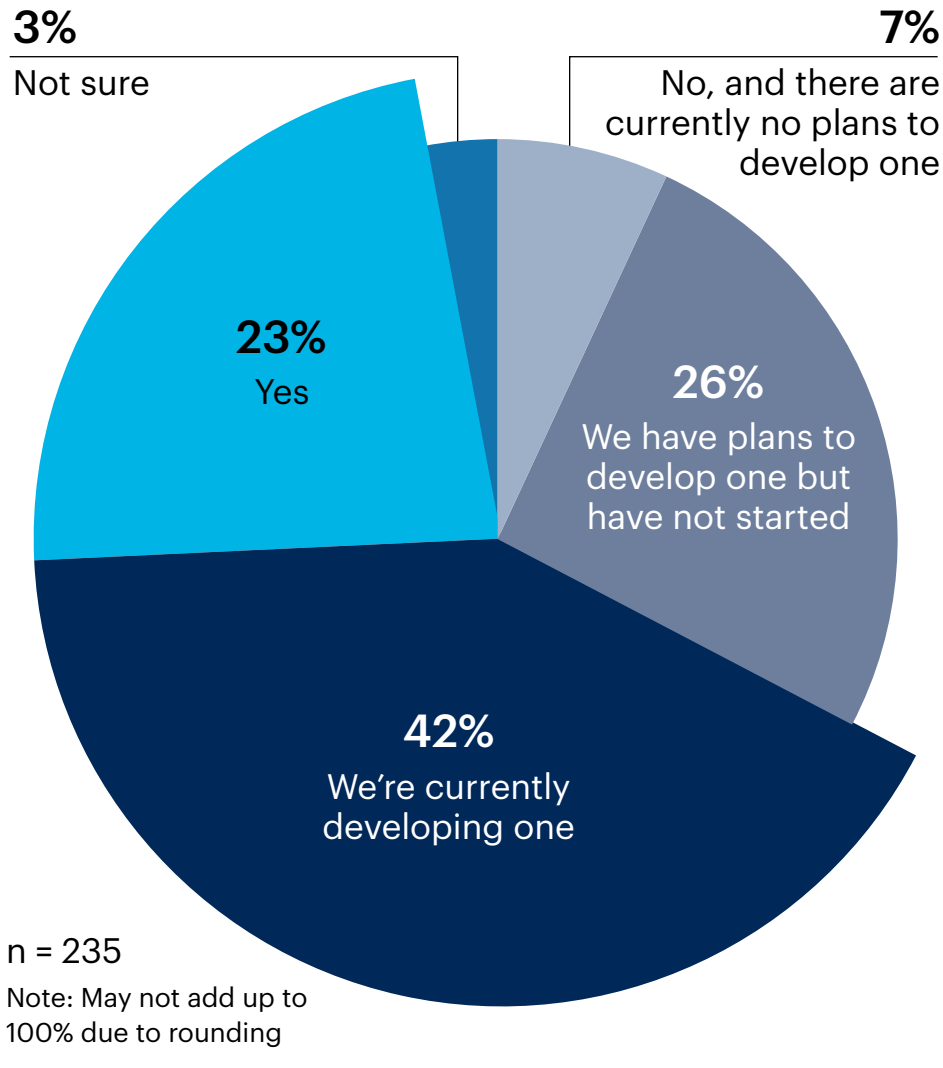

65% of all respondents (n = 235) say their organization has either established a **governance strategy for generative AI tools** or is in the process of developing one.

Has your organization established a specific governance strategy for generative AI tools?*

Full question text: With AI governance strategy being defined as, "guidelines and protocols to responsibly manage, monitor, and regulate the development, deployment, and usage of AI systems to ensure ethical, legal, and accountable outcomes," has your organization established a specific governance strategy for generative AI tools?

Over half (**56%**) are facing **challenges with team or skills gaps** when it comes to generative AI governance.

And many respondents say their organization is struggling with the regulatory landscape (46%) or unclear governance roles and responsibilities (44%).

What challenges is your organization currently facing regarding generative AI governance? Select all that apply.

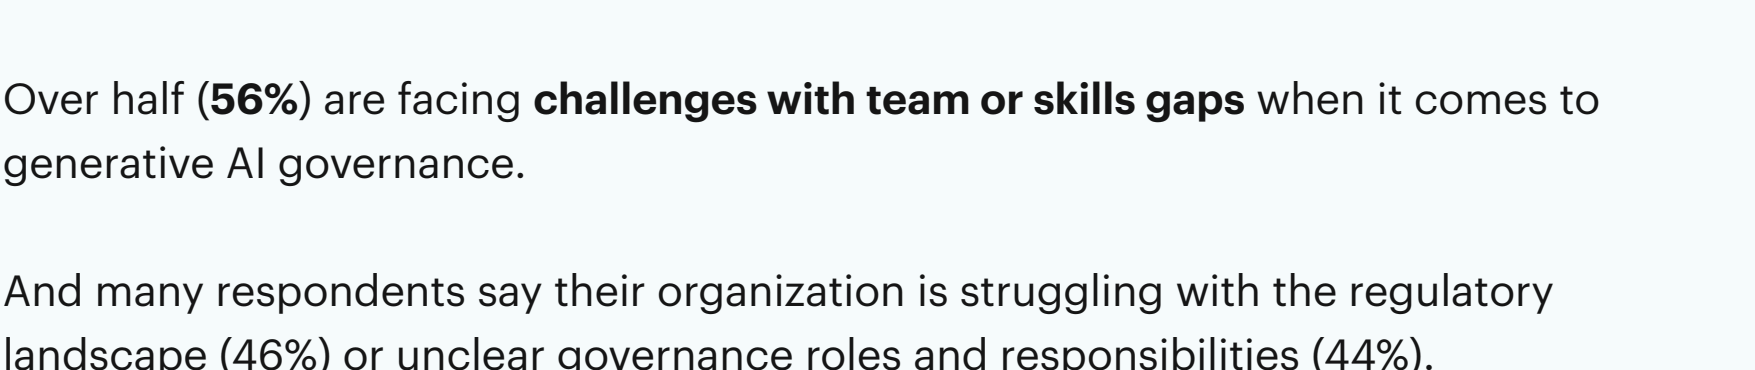

Lack of consensus on governance policies **21%** |
Decision-making power lies outside of IT/security functions **20%** |
We're not facing any challenges with generative AI governance **17%** |
Existing data governance issues **8%** | Challenge(s) not listed here **7%** |
Not sure **1%** | Other **0%**

n = 235

> "The biggest issue for us at the moment is around ensuring the output of generative AI is properly tested at scale to ensure errors are below an acceptable rate, or (for different use cases) that output is properly reviewed by a qualified human rather than a cursory review."
>
> \- C-suite, professional services industry, <1,000 employees

> "Shadow use of GenAI continues to be the biggest governance challenge. I expect tools to evolve relatively quickly to identify and block."
>
> \- VP, manufacturing industry, 10,000+ employees

Question: Please share any final thoughts on the current state of generative AI governance at your organization. Feel free to elaborate on how you think it could evolve over the next year.

Most are working on corporate policies for GenAI but over one-third have no plans to offer training on proper usage

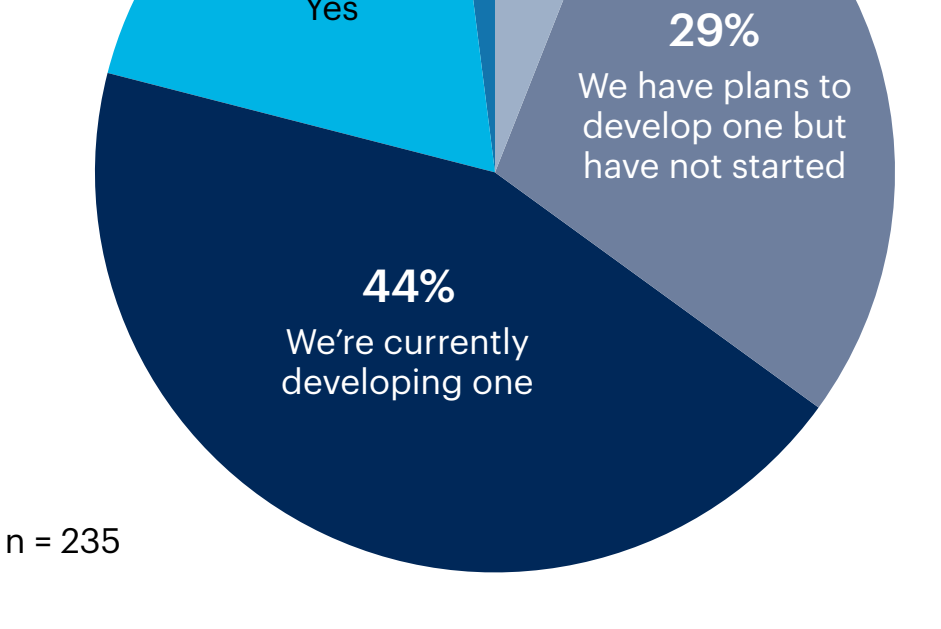

Although only **19%** of respondent organizations **currently have a corporate policy for generative AI tools,** almost three-quarters (73%) are either developing one or plan to.

Has your organization developed a corporate policy for generative AI tools?*

Full question text: With corporate policy being defined as, "a set of formal guidelines and rules adopted by a company to govern various aspects of its operations, conduct, and decision-making processes in order to ensure consistency, compliance, and alignment with its objectives," has your organization developed a corporate policy for generative AI tools?

23% of surveyed leaders say **training on proper usage of generative AI is required** at their organization, with 11% indicating that this requirement applies to all employees.

More than one-third (36%) note that training on proper generative AI usage is in development, but nearly as many (35%) report that their organization does not currently plan to offer any such training.

Are employees at your organization required to complete training on proper usage of generative AI?

> "I'm [in] the education sector and there is a lot of discussion about plagiarism but very little [that] I've seen about AI generated content or misinformation being included, unknowingly, by students or professors in citations or references."
>
> \- Director, software industry, 1,000 - 5,000 employees

> "It's a race against time. We don't want to be the ones holding things up so between now and having better governance, education is the way we're approaching it."
>
> \- Director, consumer goods industry, 1,000 - 5,000 employees

Question: Please share any final thoughts on the current state of generative AI governance at your organization. Feel free to elaborate on how you think it could evolve over the next year.

The majority are considering GenAI use cases for security, with vulnerability management as the most commonly cited application

93% of surveyed leaders with significant involvement in GenAI security or risk management efforts (n = 226) indicate that their organization is **considering generative AI use cases for security operations.**

Over half (58%) are still exploring potential use cases, but 35% report that their organization has already identified which ones to pursue.

Has your organization identified generative AI use cases to pursue for security operations?

Question shown only to respondents who answered "I own responsibility for generative AI security and/or risk management," "I am heavily involved" or "I am somewhat involved (e.g., working with other functional leaders on generative AI security and/or risk management)" to "Are you involved in security and/or risk management efforts related to the use of generative AI tools in your organization?"

n = 226

Among respondent organizations that are exploring or pursuing generative AI use cases for security operations (n = 210), the most commonly considered applications are **vulnerability management (50%), chatbots (47%)** and **reporting and/or documentation (46%).**

Over one-third list secure application development assistants (38%) or threat intelligence (37%) as use cases under consideration.

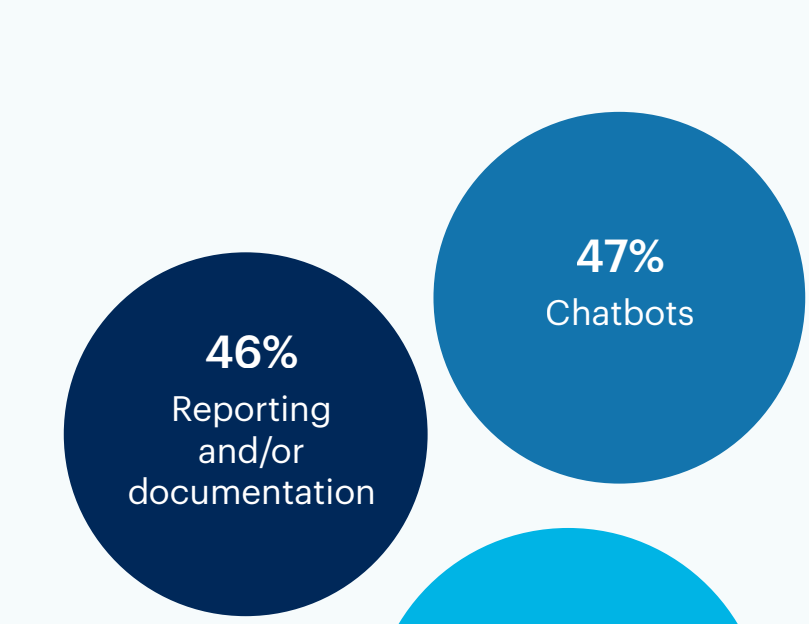

What generative AI use cases is your organization pursuing or exploring for security operations? Select all that apply.

Policy management (including generation) **29%** |
False positive reduction **28%** |
GenAI tools and services management **21%** |
Real-time risk assessment and quantification **19%** |
Zero trust (e.g., automated review of access requests) **19%** |
Asset inventory management **15%** |
Supplement security talent **11%** | Training junior security staff **8%** |
Use case(s) not listed here **5%** | Not sure **1%** | Other* **<1%**

*Other includes: SIEM / Hunting supplement

Question shown only to respondents who answered "Yes" or "We're still exploring potential use cases" to "Has your organization identified generative AI use cases to pursue for security operations?"

> "This is an exciting field. I see generative AI showing up in new ways that will continue to support more and more new ideas and uses. However, I am very skeptical of the output and am curious how this will impact adoption. For example, a false positive AI detector that has false positives would be very bad...and that is the current state of AI."
>
> \- C-suite, healthcare industry, 10,000+ employees

> "Generative AI might be the future but I still have my reservations. Vendors will give assurances and licenses but [it's] not IF something goes wrong, to me, it's a matter of WHEN. I fear all what they are offering right now will go down the drain. A bit pessimistic? YES, it is."
>
> \- C-suite, educational services industry, <1,000 employees

Question: Please share any final thoughts on the current state of generative AI governance at your organization. Feel free to elaborate on how you think it could evolve over the next year.

Model transparency and vendor privacy policies are common focuses of respondents' GenAI vendor strategies

58% of respondents with significant involvement in GenAI security or risk management efforts (n = 226) say their organization's vendor strategies include **evaluating third-party applications for model transparency requirements** or evaluating vendor privacy policies.

Over half (55%) require, or plan to require, that vendors include data governance assurances in their license agreements.

What vendor strategies is your organization using or planning to use in order to mitigate risks related to generative AI (GenAI) tools? Select all that apply.

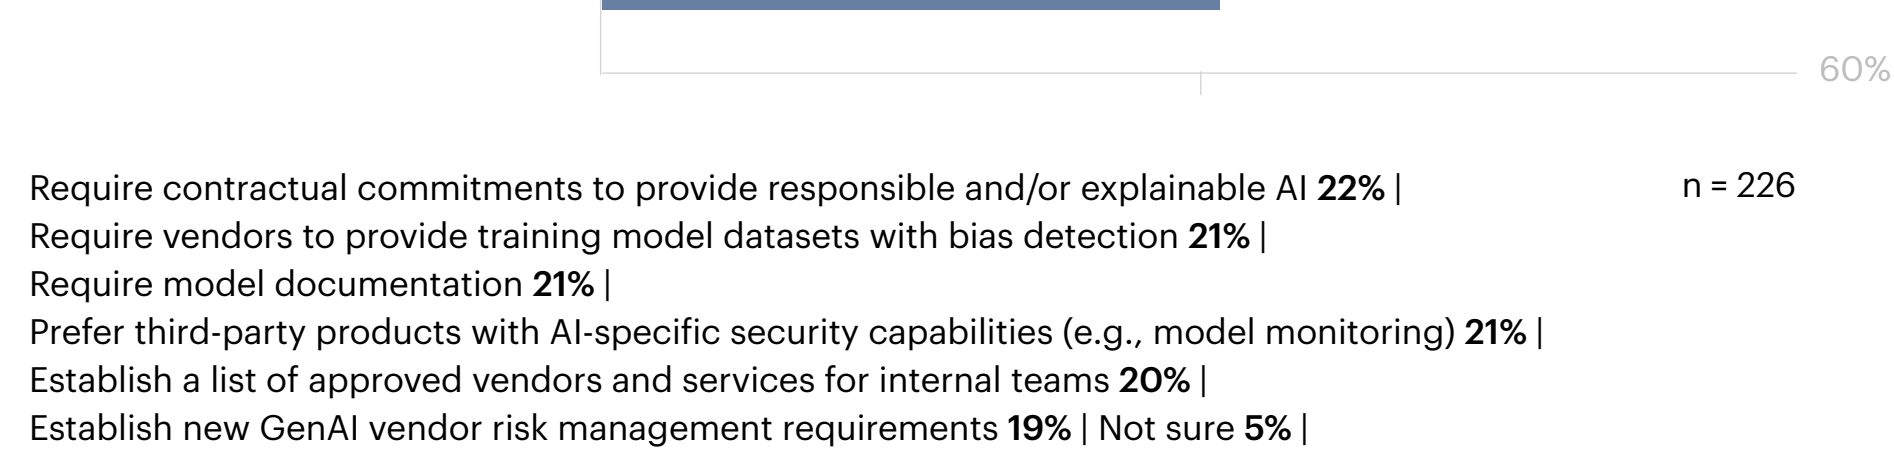

Require contractual commitments to provide responsible and/or explainable AI **22%** |
Require vendors to provide training model datasets with bias detection **21%** |
Require model documentation **21%** |
Prefer third-party products with AI-specific security capabilities (e.g., model monitoring) **21%** |
Establish a list of approved vendors and services for internal teams **20%** |
Establish new GenAI vendor risk management requirements **19%** | Not sure **5%** |
Vendor strategy(ies) not listed here **4%** | Other **0%**

n = 226

Question shown only to respondents who answered "I own responsibility for generative AI security and/or risk management," "I am heavily involved" or "I am somewhat involved (e.g., working with other functional leaders on generative AI security and/or risk management)" to "Are you involved in security and/or risk management efforts related to the use of generative AI tools in your organization?"

More than three-quarters **(76%) plan to hire staff over the next 12 months** to address generative AI governance or security needs.

One-third (**33%**) intend to hire **only internal resources** for this purpose, while 24% plan to hire both internally and through outsourcing or consulting.

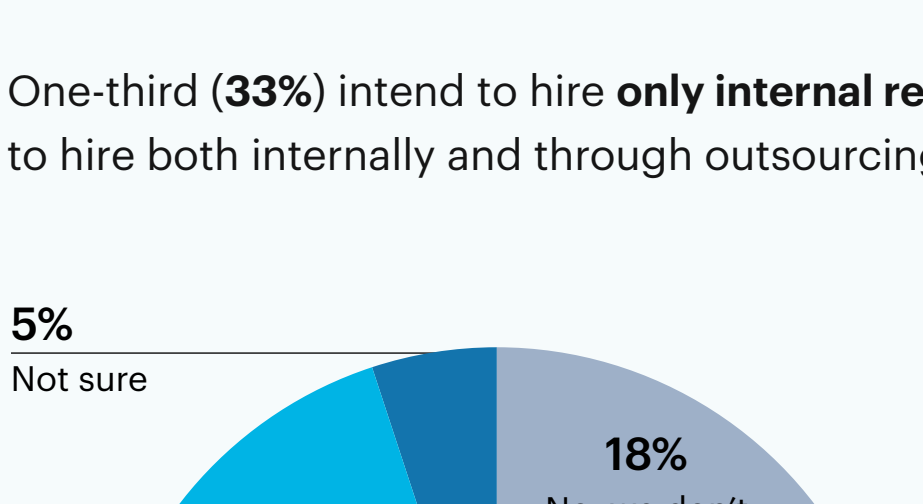

Within the next 12 months, is your organization planning to hire any staff specifically to address security and/or governance needs related to generative AI?

Note: May not add up to 100% due to rounding

Question shown only to respondents who answered "I own responsibility for generative AI security and/or risk management," "I am heavily involved" or "I am somewhat involved (e.g., working with other functional leaders on generative AI security and/or risk management)" to "Are you involved in security and/or risk management efforts related to the use of generative AI tools in your organization?"

> "We are steadily drinking the AI Kool-Aid. We have a first draft of policies that range from must be, to it would be nice, and right now the vendors we are in discussion with are hitting 95% for those."
>
> \- C-suite, utilities industry, 1,000 - 5,000 employees

> "When we use GenAI, we still review it based on [the] ISO 27001 framework."
>
> \- Director, manufacturing industry, 5,000 - 10,000 employees

Question: Please share any final thoughts on the current state of generative AI governance at your organization. Feel free to elaborate on how you think it could evolve over the next year.

Respondent Breakdown

Region

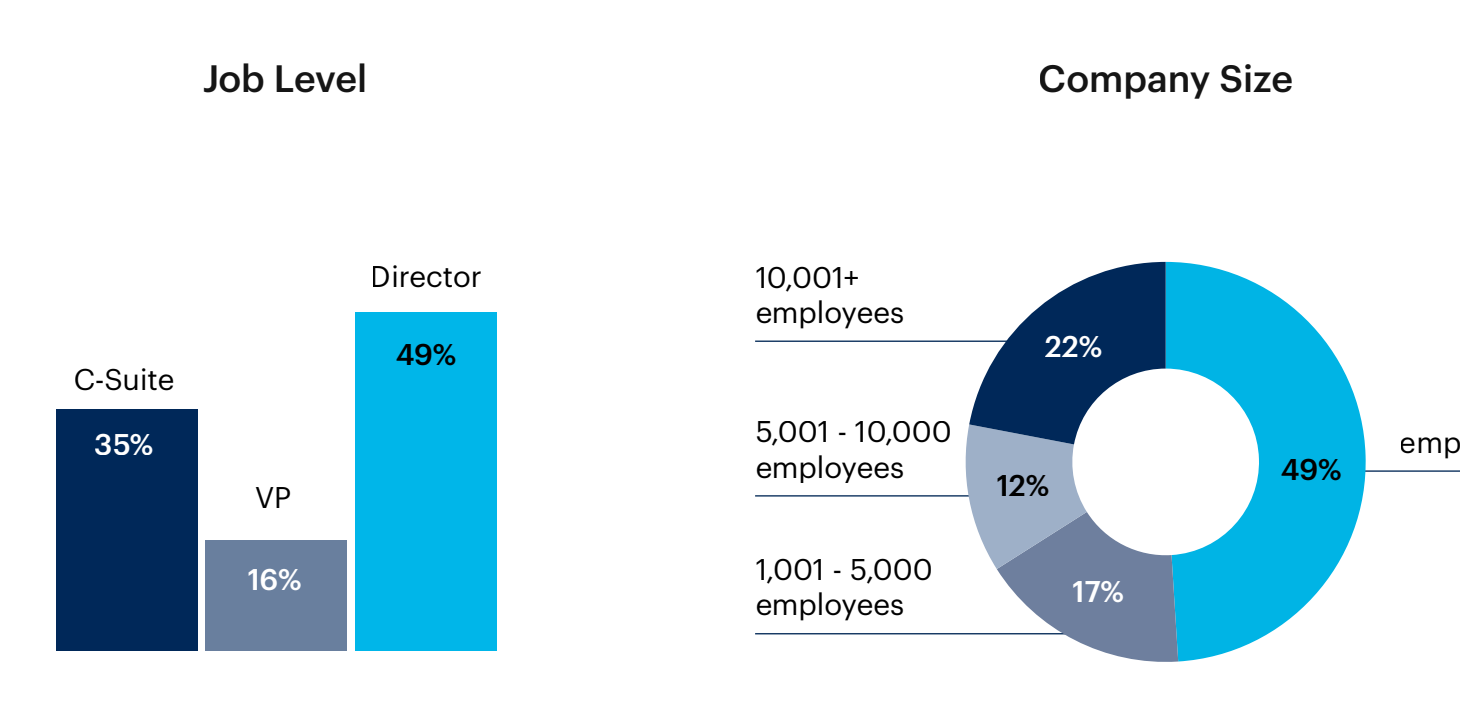

Job Level

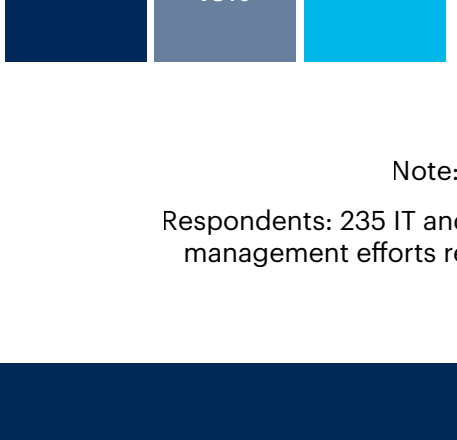

Company Size

Note: May not add up to 100% due to rounding.

Respondents: 235 IT and information security leaders involved in security and/or risk management efforts related to the use of generative AI tools at their organization

Gartner

Source: Gartner Peer Community, Navigating Generative AI Governance: Infosec Leader Perspectives survey