



Research



Shaping the Future of Generative AI

The Impact of Open Source Innovation

Adrienn Lawson, *The Linux Foundation*
Stephen Hendrick, *The Linux Foundation*
Nancy Rausch, *The Linux Foundation*
Jeffrey Sica, *The Linux Foundation*
Marco Gerosa, *Ph.D., Northern Arizona University*

Foreword by
Hilary Carter, *The Linux Foundation*

November 2024

Shaping the Future of Generative AI

84% of organizations have moderate, high, or very high **adoption of GenAI.**



For 92% of surveyed companies, **GenAI is important**, and 51% consider it extremely important.

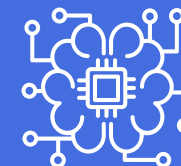


41% of GenAI infrastructure code **is open source.**



For 71% of organizations, the **open source nature of a model / tool has a positive influence** on its adoption, due to transparency and cost efficiency.

78% of organizations believe it is important to use open source tools hosted by a **neutral party**, primarily due to standards & regulations compliance and trust.



82% of respondents agree that **open source AI is critical for a positive AI future.**

30% of organizations **use proprietary data** for their proprietary models, and 22% use it for open source models.



Most organizations adopt multiple strategies for hosting GenAI inference, including **self-hosting in the cloud (49%)** and managed API services (47%).

Among those who serve or self-host GenAI models, **50% use Kubernetes** for their inference workloads.



65% of surveyed organizations **build and train GenAI models on cloud-based infrastructure.**



GenAI has **improved productivity for 79% of respondents** and has allowed them to learn new skills and improve creativity and innovation.



For the future of GenAI, 83% of respondents agree that **AI needs to be increasingly open.**



Contents

- Foreword 04
- Executive summary 05
- Introduction 07
- GenAI adoption and use in organizations 08
 - GenAI adoption 08
 - GenAI activity breakdown: Consumption dominates as custom model building gains traction 09
 - Primary GenAI use cases 12
- How open source is expanding the role of GenAI 16
 - High adopters of GenAI are more likely to use open source tools than low adopters 16
 - The critical role of open source tools and frameworks in model building and inference 18
 - The open source nature of a tool has a positive influence on its adoption 20
- GenAI and the cloud native approach 23
 - Cloud native and hybrid cloud strategies are foundational to how organizations deploy and host their GenAI models 23
 - Kubernetes as a key enabler for hosting scalable GenAI workloads 25
 - Cloud-based infrastructure leads the way in GenAI model building, with hybrid and on-premises solutions remaining key 27
- Challenges in GenAI adoption 29
 - Primary concerns of GenAI adoption 30

- Investment in GenAI 31
- Impact on employment 33
- The future of GenAI is open 34
 - Top priorities for open source projects 34
 - The role of open source AI in the future of GenAI 36
- Conclusions and recommendations 38
- Methodology 40
 - About the survey 40
 - Data.World access 42
 - Respondent demographics 42
 - About the authors 43
 - Acknowledgments 44

Foreword

A few days before this report was published, my son, who is pursuing his Bachelor of Music degree, called me to ask what I thought the impact would be of “open source Generative AI” on the music industry. “Did I think that open source Generative AI would help creators, or hurt them?” he asked. I nearly dropped the phone. Of course I took him through my reasoning for why openness was the way to build a trusted future for digital creations of all kinds, be they digital music, or digital applications used in an industry context. I even showed him some data!

We know that GenAI is transforming industries at an unprecedented pace. As this technology moves into the mainstream, organizations are rallying around the idea that AI’s future must be open. In fact, 82% of organizations believe open source AI is critical for ensuring a positive AI future, and 83% agree that AI needs to be increasingly open to foster trust, collaboration, and innovation. To me, this is the headline takeaway from this report.

The Linux Foundation is proud to champion this vision by nurturing an ecosystem where openness drives progress. Projects like PyTorch and initiatives such as the Generative AI Commons exemplify how open source fuels innovation. Meanwhile, the LF AI & Data Foundation’s Model Openness Framework and its companion tools are empowering model creators and users with practical, transparent guidance for building and adopting open AI systems.

Cloud native technologies are also central to this evolution. Not only can cloud native provide a scalable and reliable platform for running AI workloads on cloud infrastructure, but AI is enhancing cloud native offerings themselves. Through shared standards, robust frameworks, and secure infrastructure, the Cloud Native Computing Foundation (CNCF) is enabling enterprises to reduce costs and accelerate the performance of AI applications. This symbiosis underscores the transformative power of open source to meet today’s business and technical challenges.

Generative AI’s potential is limitless, but its success relies on trust, accessibility, and global collaboration. I am grateful to LF AI & Data and CNCF for sponsoring his research, and in doing so, creating data that can help decision making to scale and sustain open source AI projects.

This report is a testament to what we can achieve when the world works together, openly and transparently. For next generation creators like my son, and business decision makers, it provides a reason to be optimistic.

HILARY CARTER

Senior Vice President, Research
The Linux Foundation

Executive summary

The report, *Shaping the Future of Generative AI*, written by the Linux Foundation, supports the important role of open source in the evolution and integration of generative AI (GenAI) technologies within organizations. Based on a survey of 316 professionals across diverse industries, the report shows how open source platforms and tools are not only accelerating GenAI adoption but are also setting a foundational framework for future AI advancements. Currently, 94% of organizations are using GenAI. Leading use cases include process automation, content generation, and code generation.

Open source software is already a force shaping GenAI. On average, 41% of an organization's code infrastructure that supports GenAI is open source. Higher adopters of GenAI are more reliant on open source code (47%) compared to lower adopters (35%). Organizations that are higher adopters of GenAI are not just heavy users of open source technology; 63% are also significant contributors to open source. Consequently, 71% of respondents report that open source positively influences their decision-making, and 73% of organizations expect to increase their use of open source GenAI tools over the next two years.

Central to the success of the GenAI space are open source frameworks such as TensorFlow and PyTorch for building and training GenAI models and application frameworks including LangChain and LlamaIndex for inferencing. These open source frameworks enable organizations to build, train, and deploy models at a fraction of the cost associated with proprietary tools. Open source models empower organizations to develop customized solutions while preserving transparency and reducing dependency on closed-source platforms. This flexibility has proved essential in industries where trust, transparency, and regulatory compliance are critical.

Looking to the future of GenAI technology, open source's influence in the AI domain is expected to expand further. This survey reveals that 83% of organizations strongly agree or agree that AI needs to be increasingly open. Additionally, 82% report that open source AI is a critical component for a sustainable AI future, with 61% expressing confidence that the benefits of open source outweigh the associated risks. The growth of open source GenAI technology is likely to be significant, with 73% of organizations expecting to increase their use of open source generative tools over the next two years and 26% anticipating a substantial rise in use.

Organizations that integrate open source GenAI tools not only benefit from reduced costs but also often contribute to a collaborative ecosystem that drives technological progress. The report also discusses cloud native's critical role in supporting scalable GenAI solutions. Cloud-based infrastructure, combined with open source frameworks and tools, allows organizations to manage and deploy complex AI models more efficiently. Kubernetes, for instance, has emerged as a key enabler for orchestrating scalable GenAI workloads, with 50% of organizations using Kubernetes to host some or all of their GenAI inferencing workloads.

“Organizations with higher levels of GenAI adoption are helping to shape next-generation frameworks and models, aligning them more closely with advanced, real-world use cases.”

This report recommends that organizations continue to prioritize open source in their GenAI strategies to remain competitive and aligned with industry trends. It also highlights the importance of neutral organizations, such as the Linux Foundation, Cloud Native Computing Foundation (CNCF), and LF AI & Data Foundation, in providing open governance structures that improve trust and collaboration. As AI continues

to reshape industries, open source will remain indispensable, offering a balanced, transparent, and community-led pathway to innovation that will define the future of AI technologies. By offering accessible, adaptable, and community-driven resources, open source has democratized access to GenAI, allowing organizations of all sizes to leverage cutting-edge AI capabilities securely and effectively.



Introduction

This report explores the deployment, use, and challenges of GenAI technologies in organizations and the role and future of open source in this domain.

Linux Foundation Research and its partners conducted a web survey from August through September 2024, which provided the empirical basis for this study. Survey respondent screening ensured that respondents:

- Were familiar, very familiar, or extremely familiar with the adoption of GenAI in their organization

- Worked for an organization
- Had professional experience

A total of 316 respondents completed the survey.

There are also a variety of callouts throughout this report. These callouts include selected verbatim comments in (italicized blue text with no background colour) response to an open text question in the survey, which asked, “Do you have any final comments or thoughts about GenAI?”



GenAI adoption and use in organizations

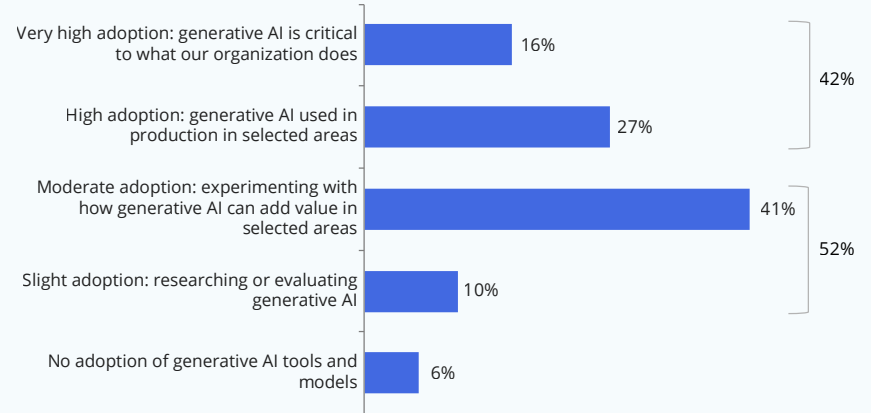
Organizations are adopting GenAI because of its ability to address a broad array of strategic and tactical needs, including content creation, personalized customer experiences, decision support, process automation, employee training, and research and planning. To understand the development and use of GenAI and how open source is impacting the evolution of GenAI, we need to first evaluate how organizations are involved with GenAI, its leading use cases, and the maturity of GenAI deployments.

GenAI adoption

Figure 1 shows the extent of the adoption of GenAI. The top chart in Figure 1 indicates that 94% of organizations are involved with GenAI and can be segmented into two categories: organizations who have very high or high GenAI adoption (42%, higher GenAI adopters) and organizations who have slight or moderate GenAI adoption (52%, lower GenAI adopters). We also see in Figure 1 that 84% of organizations have a moderate, high, or very high adoption of GenAI.

FIGURE 1: HOW ORGANIZATIONS ARE ADOPTING GEN AI

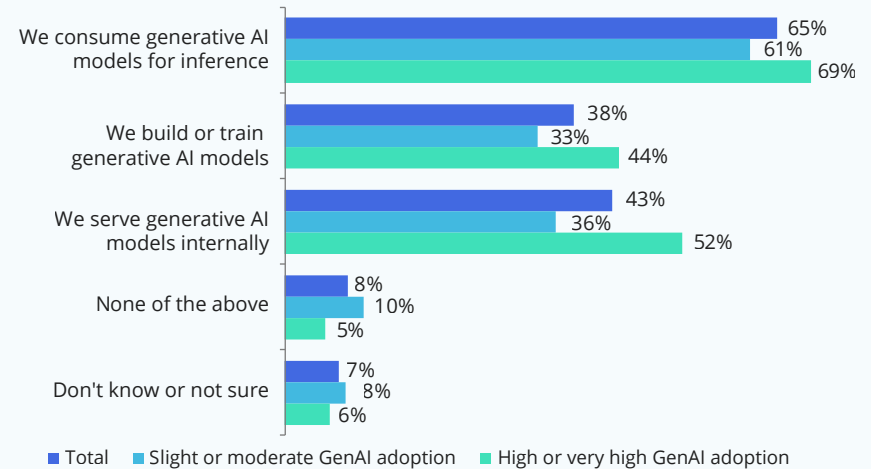
To what extent has your organization adopt generative AI? (select one)



2024 GenAI survey, Q7, Sample Size = 316

What activities does your organization undertake with generative AI models? (select all that apply) segmented by:

To what extent has your organization adopted GenAI? (select one)



2024 GenAI survey, Q32 by Q7, Sample Size = 297, Valid Cases = 297, Total Mentions = 479, answered by organizations who adopted GenAI in Q7

GenAI activity breakdown: Consumption dominates as custom model building gains traction

Core activities related to GenAI, including building (training), serving, and inferencing (consuming a model), are shown in the chart at the bottom of Figure 1. Inferencing, at 65% overall, is a primary GenAI activity. Inferencing is significantly higher than either building models (38%) or serving these models (43%). Organizations are choosing to tune and / or train their own GenAI models to meet specific business needs and make these models more accurate and relevant. Custom models allow organizations to tailor responses, fine-tune language, and incorporate domain-specific knowledge to create outputs that align closely with their brand and industry requirements. Fine-tuning a model also provides enhanced control over the AI's evolution and reduces dependency on external providers. Figure 1 (at the bottom) also shows that organizations that have a higher level of GenAI adoption also are more involved in building / training models (44%), internally serving these models (52%), and consuming these models (69%).

Figure 2 explores the various techniques in use to improve the performance of GenAI models. The leading technique, prompt engineering, is showing significant gains for nearly 80% of organizations that have adopted GenAI. Prompt engineering is the practice of optimizing inputs (prompts) to deliver the

most accurate, relevant, or creative outputs from an AI model. By carefully designing prompts, prompt engineers improve model performance. The elegance of prompt engineering is that this increased performance does not require any changes to the underlying GenAI model, although it does require a more detailed approach in defining inputs.

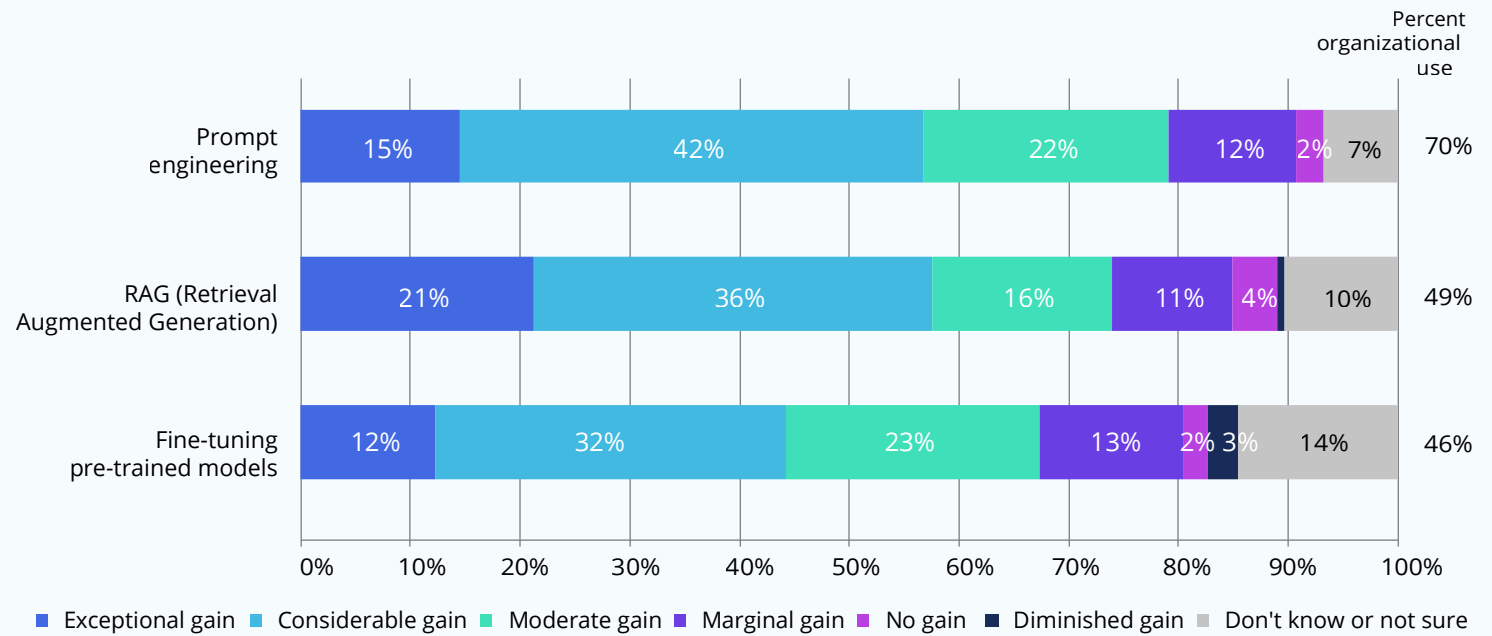
Retrieval augmentation generation (RAG) is also a leading technique for improving performance. RAG combines the power of large language models (LLMs) with real-time, relevant information retrieval to generate highly informed and contextually accurate responses. This approach augments the model's outputs by grounding them in domain-specific information. RAG improves model performance by providing a bridge between static model knowledge and dynamic, up-to-date content, which is ideal for applications such as customer support, research, and decision support systems. RAG is driving material gains for more than 70% of organizations that have adopted GenAI.

Fine-tuning an LLM is yet another technique that organizations commonly use to improve the performance of their GenAI models. Fine-tuning adjusts the LLM's internal parameters by training on domain-specific data, which embeds specialized knowledge directly into the model. This makes the model more fluent in specific topics but limits it to static knowledge present during training. Fine-tuning is showing significant gains for nearly 70% of organizations using GenAI.



FIGURE 2: THE TOP THREE TECHNIQUES FOR IMPROVING GEN AI MODEL PERFORMANCE

How much have the following generative AI techniques improved the performance of your baseline approach?
 (one response per row) filtered for: **What techniques are you using in your organization?** (top three shown)



2024 GenAI survey, Q33, Sample Size = 297, Valid Cases = 297, Total Mentions = 905, answered by organizations who adopted GenAI in Q7
 2024 GenAI survey, Q34, Sample size = 206 to 138, sorted by the sum of "Exceptional, considerable, and moderate gain"

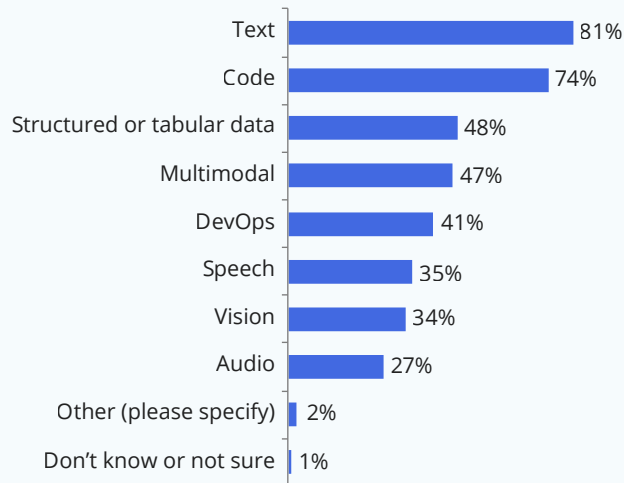


The left-hand chart in Figure 3 shows that text (81%), code (74%), and structured or tabular data (48%) are the leading GenAI modalities. Text, code, and structured data are the most common modalities for GenAI because they are widely available, interpretable, and foundational to a broad range of applications. Text data, which LLMs support, covers a wide spectrum of natural language applications, enabling models to generate coherent responses, summaries, translations, and other human

language outputs. Code, as a logical and rule-based language, is highly suited for automating tasks, generating scripts, and supporting software development. Structured data—such as tables, databases, and labeled datasets—provide organized information that GenAI can use for pattern recognition and data synthesis in areas such as decision support and recommendations.

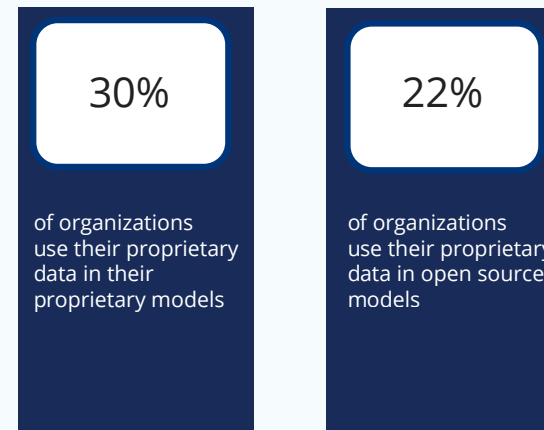
FIGURE 3: COMMON GEN AI MODALITIES AND DATA USE

What generative AI modalities are you using or planning to use in your organization? (check all that apply)



2024 GenAI survey, Q31, Sample Size = 297, Valid Cases = 297, Total Mentions = 1,165, answered by organizations who adopted GenAI in Q7

Does your organization have proprietary data that could be used to train or improve the performance of generative AI models? (select all that apply)



2024 GenAI survey, Q24, Sample Size = 297, Valid Cases = 297, Total Mentions = 473, answered by organizations who adopted GenAI in Q7

The right-hand chart in Figure 3 shows the percentage of organizations using their proprietary data to improve the performance of their proprietary GenAI model (30%) or open source GenAI model (22%). Some organizations use their own data to train both proprietary and open source models. When we redistribute the data based on these three categories, we

find that organizations use proprietary data to improve the performance in 22% of proprietary models, 13% of open source models, and 9% with both models. This yields a total of 44% of organizations that are using proprietary data to improve their models.

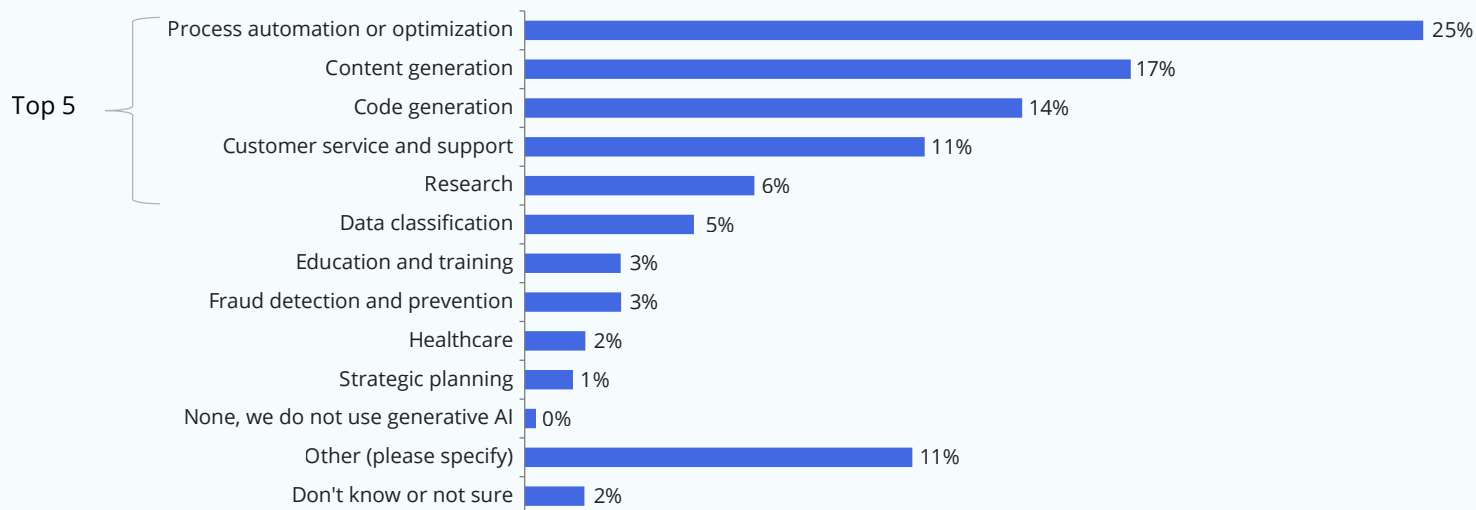
Primary GenAI use cases

Organizations are using GenAI in many ways, although there are five primary use cases. Figure 4 shows that the leading primary use case is process optimization or automation (25%) followed by content generation (17%), code generation (14%), customer service and support (11%), and research (6%).

Process automation / optimization is the leading GenAI use case because it offers businesses transformative efficiency, reduces manual tasks and errors, and decreases operational costs. With its capacity for natural language understanding and response, GenAI can handle diverse queries and tasks, providing a more adaptable and scalable approach to automation. By identifying patterns and recommending improvements, GenAI not only streamlines processes but also creates room for innovation.

FIGURE 4: PRIMARY GEN AI USE CASES

What's your organization's primary use case for generative AI? (select one)



2024 GenAI survey, Q13, Sample Size = 297, answered by organizations who adopted GenAI in Q7

GenAI is also useful for content (17%) and code (14%) generation. For content creation, GenAI can generate articles, blogs, and marketing materials in seconds, reducing the workload and ensuring consistency in style and tone. It enhances ideation, delivering varied perspectives or outlines that help teams

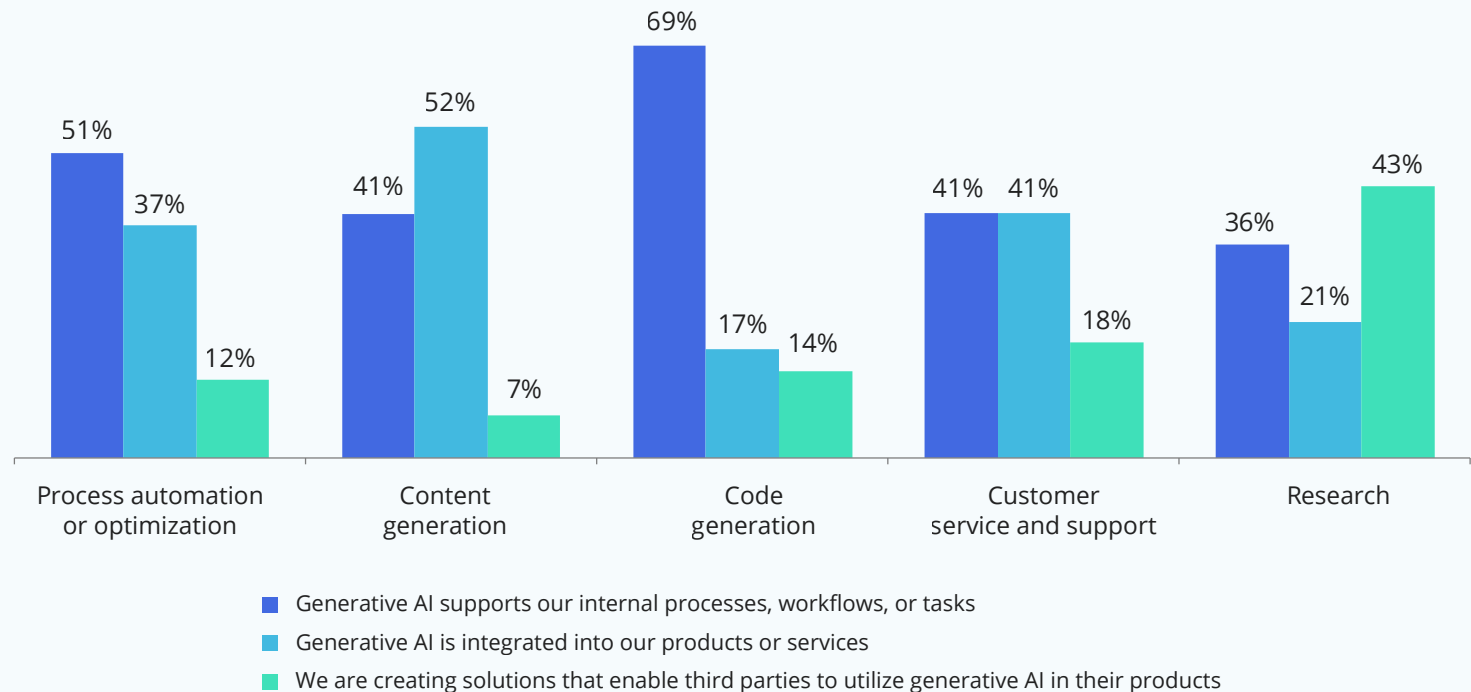
focus on refinement. For code generation, GenAI accelerates development, offering quick prototypes, code suggestions, and debugging support. By automating routine coding tasks, it reduces errors and frees developers to focus on complex problem solving.

Gen AI can support customer service (11%) by providing instant, constant support through intelligent chatbots and virtual assistants. It can improve response times, handle large volumes of queries simultaneously, and provide accurate, context-aware answers. GenAI can also personalize interactions by analyzing customer history and preferences, delivering tailored solutions and proactive recommendations. Additionally, it automates repetitive inquiries, allowing human agents to focus on complex cases that require a personal touch.

Figure 5 shows the five leading use cases in Figure 4 segmented by how the use case is integrated into the organization's business. Figure 5 shows that each of these five primary use cases is uniquely integrated into the organizations that identified it.

FIGURE 5: PRIMARY GEN AI USE CASES SEGMENTED BY INTEGRATION INTO THE BUSINESS

What's your organization's primary use case for generative AI? (select one) segmented by **How is your primary generative AI use case integrated into your business?** (select one)



2024 GenAI survey, Q13 top 5 by Q15, Sample Size = 166, DKNS and Too soon to tell responses excluded from the analysis

Process automation or optimization shows a relatively high level of support for internal processes (51%) but a lesser degree of integration with an organization's products and services (37%). This lesser degree of integration reflects the complexity of aligning Gen AI models with organizational workflows

The integration of content generation (52%), in contrast, is more readily achievable, because an organization's content is already in a highly consumable form for GenAI. Code generation sees a significantly high level of support for internal activities (69%) in part because code generation provides a bounded domain that offers useful results without an excessive degree of integration or customization to an organization's environment. However, there are several reasons why the integration of code generation is just 17%, which we'll explore in Figure 6.

Customer service has a relatively high level of support for internal processes (41%) as well as integration with products and services (41%). The elevated interest in creating solutions by third parties (18%) reflects the fact that every organization has to staff customer service and support activities, so the payoff in developing an effective solution is considerable.

“GenAI applies a mathematical model to an inherently subjective process. It will never be good enough in a generic context because different people want different and incompatible things from it. However, it has the potential to become good enough in small, specialized uses, if it can stop hallucinating.”

Research is another use case where advanced analytics tool providers see tremendous opportunity for creating solutions. Because most data includes metadata, data is often structured, which improves its specificity while simplifying how organizations can use it. While the internal use cases for research is intuitively clear, the deployment of such systems is still in its infancy.

Figure 6 shows the five leading use cases from Figure 4 segmented by the level of adoption.

Content generation is the most highly integrated use case (Figure 4) is also shown in Figure 5 as the most highly fully deployed use case (22%) and has the highest level of initial production deployments (46%). The reason for this is likely due to a much higher ROI than other use cases due to the relative ease of implementation and the significant operational value add.

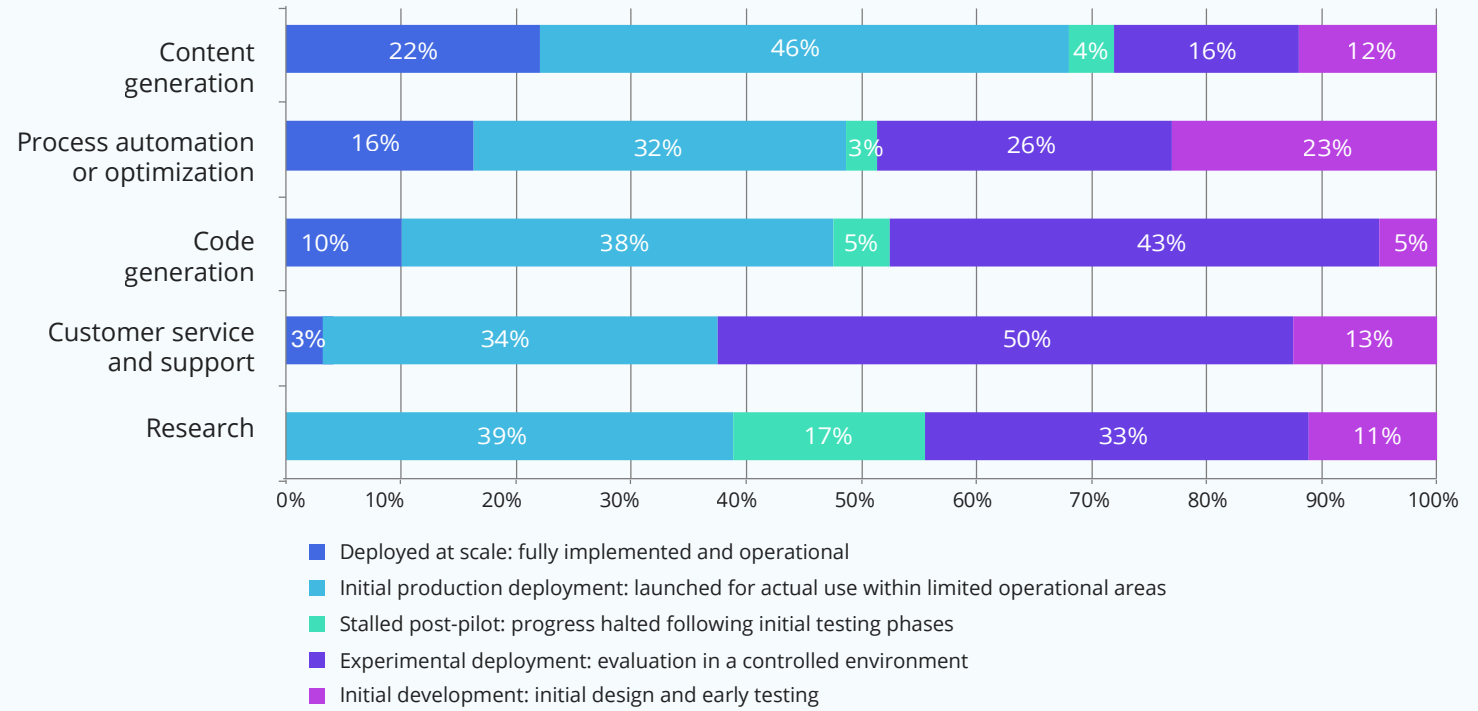
Process automation is fully deployed in just 16% of organizations using GenAI and an additional 32% are in initial production deployment. This reflects the challenges of defining the scope necessary for process automation with a GenAI model.

Both code generation (code specific GenAI models) and customer service and support (LLMs) show low full deployment rates but a very high level of initial deployment and experimental deployments suggesting that full deployment rates could increase significantly if these initial and experimental deployments show promise.

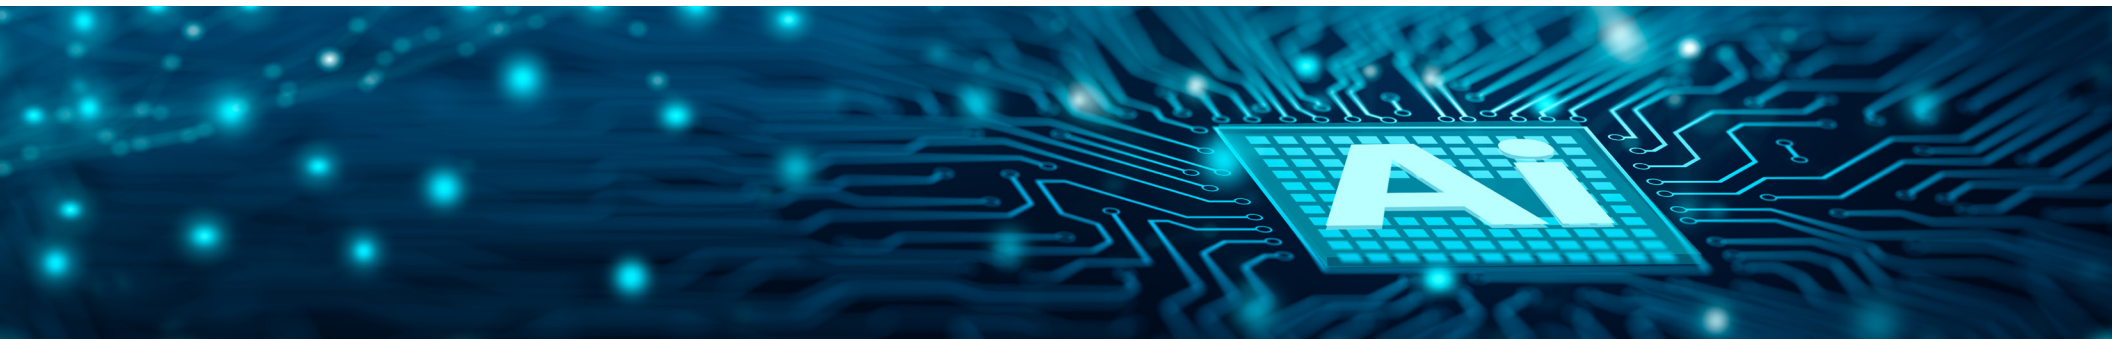
Although no large-scale deployments of the research use case are yet reported, high levels of initial deployment (39%) and experimentation (33%) are encouraging although the high level of stalled post-pilot instances (17%) is concerning.

FIGURE 6: PRIMARY GEN AI USE CASES SEGMENTED BY LEVEL OF ADOPTION

What’s your organization’s primary use case for generative AI? (select one) segmented by:
 What is the current stage of your organization’s primary generative AI use case? (select one)



2024 GenAI survey, Q13 by Q14 top 5, Sample Size = 287, DKNS responses excluded from the analysis, Q14 answered by organizations who adopted GenAI in Q7



How open source is expanding the role of GenAI

Building on the previous discussion of GenAI adoption rates, techniques, and use cases, this section explores the critical role of open source in the GenAI space. High adopters of generative tools are especially likely to rely on open source tools, which

are critical for model building, deployment, and inference. We will also look at how the open source nature of a tool positively influences its adoption.

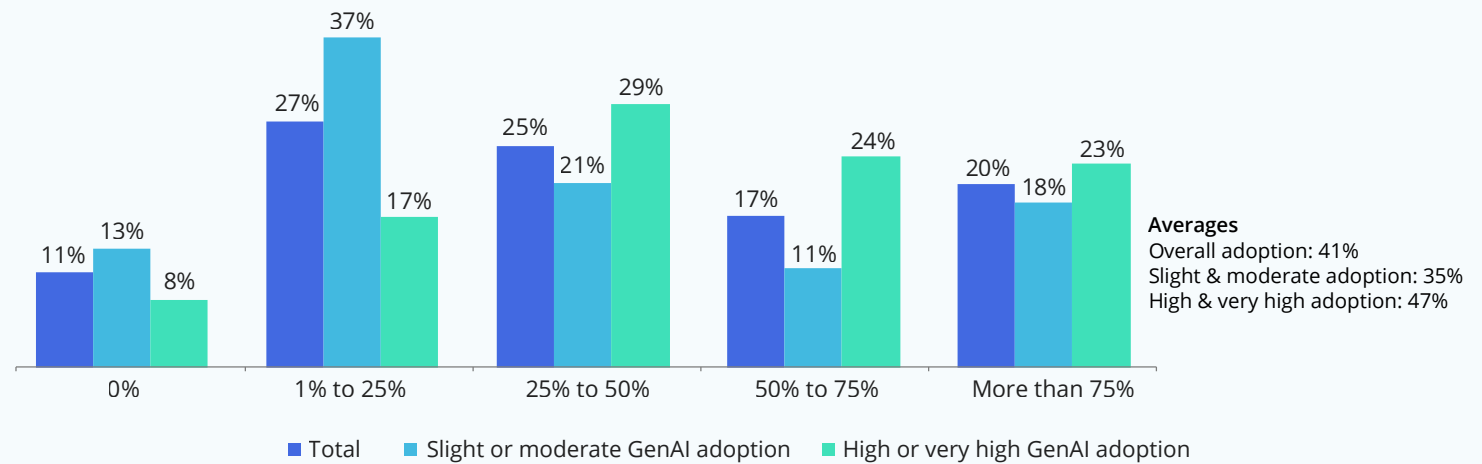
High adopters of GenAI are more likely to use open source tools than low adopters

The adoption of open source infrastructure is emerging as a strategic differentiator for organizations highly committed to GenAI. Figure 7 indicates that organizations with high or very high levels of GenAI adoption are increasingly integrating open source tools into their technology stacks, with open

source infrastructure constituting an average of 47% of their GenAI code infrastructure. In contrast, organizations with only slight or moderate GenAI usage report lower open source adoption levels, averaging 35%. Organizations with advanced AI capabilities often require the adaptability and control that open source frameworks, tools, and models provide, especially for tailoring models and pipelines to meet specialized requirements.

FIGURE 7: PERCENTAGE OF OPEN SOURCE GEN AI CODE INFRASTRUCTURE SEGMENTED BY GEN AI ADOPTION LEVEL

Approximately how much of your organization's code infrastructure that supports generative AI initiatives is currently from open source? (select one) segmented by **To what extent does your organization adopt generative AI? (select one)**



2024 GenAI survey, Q19, by Q7, Sample Size = 255, DKNS responses excluded from the analysis, answered by organizations who adopted GenAI in Q7

The implications are clear: For organizations aiming to move from experimental to production-grade GenAI, open source solutions offer a powerful foundation for both innovation and operational resilience. Open source communities continuously push the boundaries of model architectures, tools, and libraries. They can offer high-adoption organizations early access to the latest advancements in areas such as transformer models, distributed computing frameworks, and model optimization techniques.

High GenAI adopters contribute more actively to open source projects than low adopters

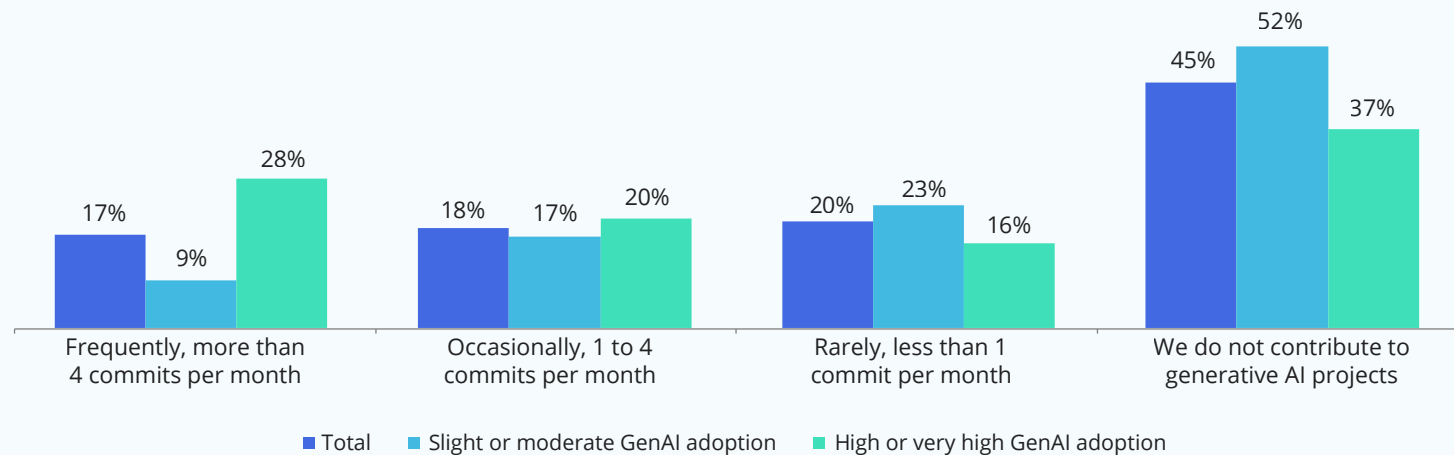
Figure 8 shows that organizations with a high adoption of GenAI are not just heavy users of open source technology; they are also significant contributors to it. High GenAI adopters are three

times more likely than their lower adoption counterparts to frequently contribute to open source GenAI projects (defined as more than four commits per month) and to open source projects related to GenAI. Specifically, 28% of high-adoption organizations report frequent contributions, compared to only 9% among slight or moderate adopters.

“28% of high-adoption organizations report frequent contributions, compared to only 9% among slight or moderate adopters.”

FIGURE 8: FREQUENCY OF OPEN SOURCE CONTRIBUTIONS SEGMENTED BY GEN AI ADOPTION LEVEL

How much does your organization contribute to open source projects related to generative AI? (select one) segmented by To what extent does your organization adopt generative AI? (select one)



2024 GenAI survey, Q21 by Q7, Sample Size = 246 (DKNS excluded), Q21 answered by organizations who adopted GenAI in Q7

“High GenAI adopters recognize the mutual benefits of engaging in the open source ecosystem.”

By contributing code and fixes, high GenAI adopters not only support the broader AI community but also gain opportunities to shape tools and frameworks to better fit their advanced use cases. Furthermore, active involvement in open source projects

The critical role of open source tools and frameworks in model building and inference

As high adopters leverage the benefits of open source to drive innovation, the open source frameworks and tools become foundational pieces in the AI/ML ecosystem. As shown in the top panel of Figure 9, PyTorch and TensorFlow, both open source, lead the model training space, with 63% and 50% adoption rates, respectively. These open source frameworks have driven major

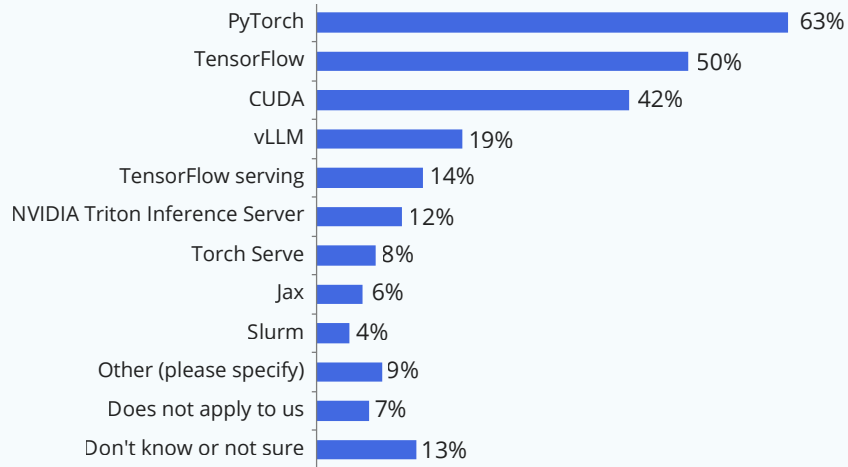
allows high adopters to stay at the forefront of AI technology and attract cutting-edge talent to their companies, especially with proprietary models being closed off, which does not allow room for inspection and experimentation from outside organizations.

advances in AI and laid the groundwork for today’s sophisticated closed-source generative models. The most common proprietary AI models are built upon the foundational research, community contributions, and innovations developed using these open platforms. Open source tools and libraries have not only accelerated research but also established standardized methods and architectures that commercial models often leverage and extend.



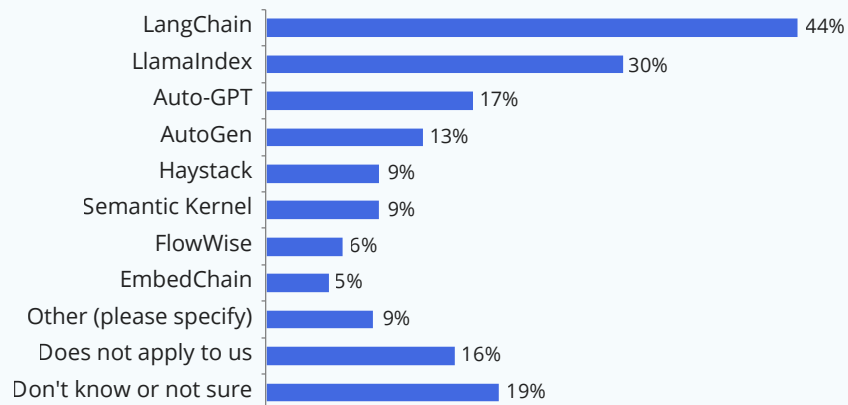
FIGURE 9: MOST POPULAR FRAMEWORKS FOR MODEL BUILDING AND INFERENCE

For building or training generative AI models, which of these frameworks does your organization currently use? (select all that apply)



2024 GenAI survey, Q37, Sample Size = 113, Valid Cases = 113, Total Mentions = 279, answered by organizations who train or build models in Q32

For model inference, which application frameworks do you use or plan to use? (select all that apply)



2024 GenAI survey, Q36, Sample Size = 193, Valid Cases = 193, Total Mentions = 341, answered by organizations who use GenAI for inference in Q32



“It’s difficult to predict where this will go, but we will certainly need new ways to obtain quality training data and to evaluate models for security and alignment. I think it’s likely that future applications will become much more complex combinations of many types of models and frameworks, and the networking of such systems will eventually lead to general AI and self-improving intelligence.”

Furthermore, the open source nature of PyTorch and TensorFlow has fostered an ecosystem where tools, pre-trained models, and best practices are freely accessible, significantly lowering the barriers to entry for organizations and researchers globally. This democratization of AI/ML development has created a collaborative environment where even proprietary models benefit indirectly from ongoing improvements and innovations within the open source community.

Similarly, for model inference, LangChain and LlamaIndex—both open source tools—are highly favored, with 44% and 30% adoption rates, respectively. These tools allow modular and adaptable solutions, which is especially critical in complex inference pipelines where organizations require tailored workflows and integrations with external data.

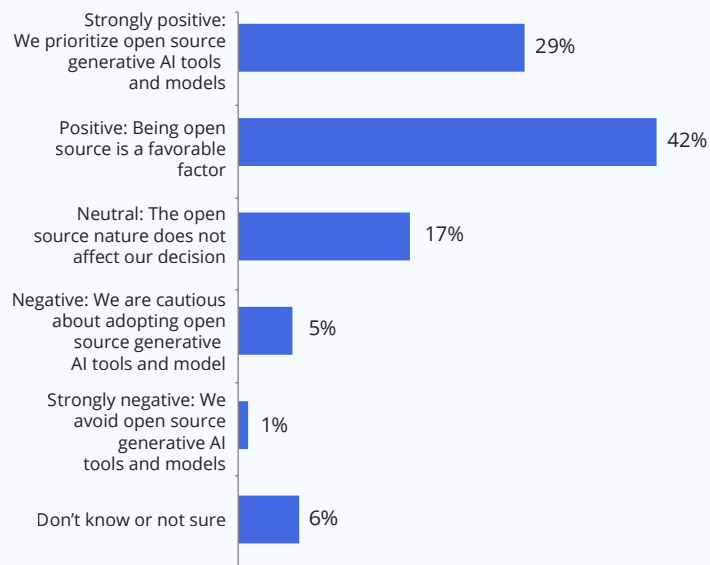
The strong reliance on open source frameworks also shows the community’s demand for interoperability and flexibility. In AI/ML, where rapid iteration and experimentation are essential, open source frameworks allow developers to integrate various tools, switch components easily, and optimize solutions across different infrastructures. This flexibility is particularly evident in the case of frameworks such as vLLM, which enables efficient LLM inference and allows the community to innovate on top of it without restrictive licensing or vendor lock-in.

The open source nature of a tool has a positive influence on its adoption

The critical role of open source frameworks and tools emphasizes the value of open source definitions that clarify what open source AI truly entails. Thanks to a lot of effort in this space, our understanding of open source AI models continues to improve. Examples include the delineation of [open source models and open science models](#) by LF AI & Data Foundation, the release of the v1 of the [Open Source AI Definition](#) (OSAID), and frameworks such as the [Model Openness Framework](#) (MOF) and its accompanying [Model Openness Tool](#) (MOT). In other areas of the AI ecosystem, open source tools—such as PyTorch, as discussed above—are already clearly defined. Figure 10 shows that a clear majority of organizations recognize open source GenAI tools and models as a positive factor in their adoption strategies. According to our survey, 71% of respondents report that open source positively influences their decision-making, with 29% prioritizing open source tools outright and 42% viewing it as a favorable factor.

FIGURE 10: IMPACT OF THE OPEN SOURCE NATURE OF A TOOL ON ITS ADOPTION

How does the open source nature of a tool or model influence its adoption within your organization? (select one)



2024 GenAI survey, Q22, Sample Size = 297, answered by organizations who adopted GenAI in Q7

Which options justify your answer to the above question? (select all that apply)

Transparency and trust in source code	52%
Cost efficiency	46%
Security	32%

2024 GenAI survey, Q23, Sample Size = 297, Valid Cases = 297, Total Mentions = 1,286, answered by organizations who adopted GenAI in Q7

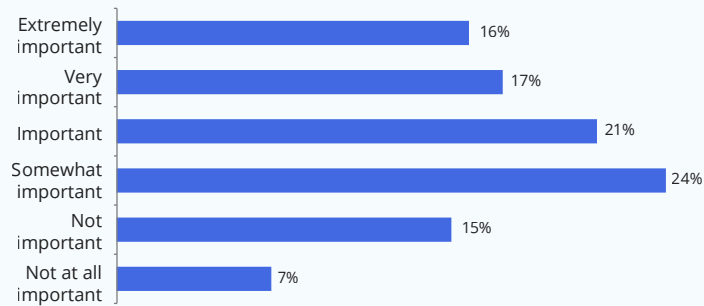
Transparency and trust of open source tools largely drive their appeal. More than half of respondents (52%) indicated that transparency and trust in the source code were the primary reasons for choosing open source GenAI solutions. Access to source code allows organizations to verify model behavior, identify potential biases, and ensure regulatory compliance—key considerations in regulated industries such as finance, healthcare, and the public sector. This transparency also supports security, as organizations can scrutinize and audit open source code to detect vulnerabilities, which is an aspect that 32% of respondents cited. Cost efficiency is another significant factor, which 46% of organizations cited. Open source tools reduce both upfront licensing costs and long-term dependence on proprietary vendors, freeing organizations to allocate resources to differentiation, customized development, and optimization.

The importance of open source tools hosted by a neutral party

Given the importance of a tool's open source nature to organizations, it is worth taking this a step further to consider the role of open governance. Foundations play a crucial role here, serving as neutral entities that can oversee and maintain open governance structures. Figure 11 shows that a majority, 78% of respondents, indicate that using open source AI tools hosted by a neutral third party is at least "somewhat important." This is broken down as follows: 16% consider *it extremely important*, 17% *very important*, 21% *important*, and 24% *somewhat important*.

FIGURE 11: IMPORTANCE OF OPEN SOURCE TOOLS THAT ARE HOSTED BY A NEUTRAL PARTY

How important is it for your organization to use AI open source tools that are hosted by a neutral party, such as the Linux Foundation, instead of a corporate entity? (select one)



2024 GenAI survey, Q25, Sample Size = 263 (DKNS excluded), answered by organizations who adopted GenAI in Q7

Open governance is needed alongside open sourcing to ensure that organizations manage projects transparently, inclusively, and with community-driven decision-making. This structure allows diverse stakeholders—developers, researchers, and end users—to have a say in how tools evolve. This autonomy helps prevent control by any single organization. For GenAI tools, where ethical concerns, security, and adaptability are paramount, open governance fosters trust and helps the technology to develop in a way that serves the wider community.



GenAI and the cloud native approach

The widespread adoption of open source tools and frameworks has set a strong foundation for innovation and collaboration in AI/ML. However, if organizations plan to integrate GenAI models further into their products or processes, they will be looking for scalable infrastructure to support model deployment and operation. This brings us to the critical role of cloud native approaches, which provide the elasticity and efficiency required to manage large-scale AI workloads.

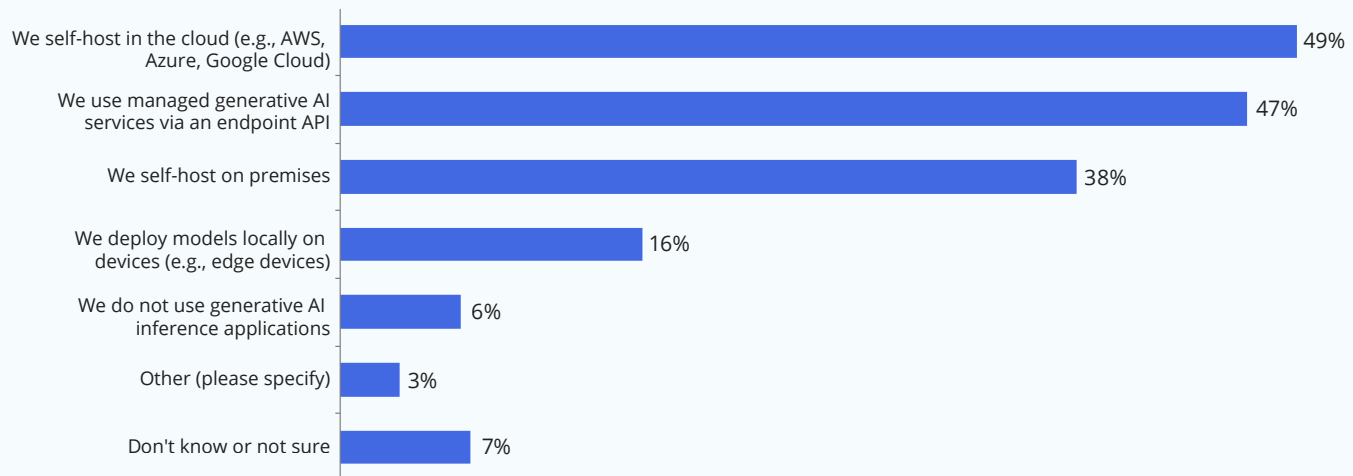
“Gen AI is a revolutionary technology. Adapting to it is the best way to survive in the competitive market.”

Cloud native and hybrid cloud strategies are foundational to how organizations deploy and host their GenAI models

Figure 12 provides insights into the hosting strategies organizations employ for deploying GenAI models and the dominance of cloud-based hosting. Specifically, 49% of respondents report self-hosting GenAI models in the cloud, while 47% leverage managed GenAI services accessible via endpoint APIs.

FIGURE 12: PRIMARY HOSTING LOCATIONS FOR GEN AI MODEL INFERENCE APPLICATIONS

Where does your organization host the generative AI models for inference applications? (select all that apply)



2024 GenAI survey, Q38, Sample Size = 193, Valid Cases = 193, Total Mentions = 319, answered by organizations who use GenAI for inference in Q32

Self-hosting in the cloud enables companies to retain greater control over their models and customize deployments to meet specific organizational needs, while still benefiting from cloud native capabilities. However, the nearly equivalent reliance on managed services (47%) reflects the value placed on reducing the operational complexity of maintaining GenAI models. Managed services via APIs also allow quicker integration into applications, minimizing the need for in-house expertise in model maintenance and optimization.

Beyond cloud hosting, 38% of organizations self-host generative models on premises. This option may appeal to organizations with strict data privacy or regulatory requirements, such as those in finance or healthcare, or those with existing on-premises infrastructure that they wish to leverage. Meanwhile, 16% of respondents deploy models locally on devices, such as edge devices. This approach is well-suited to use cases that

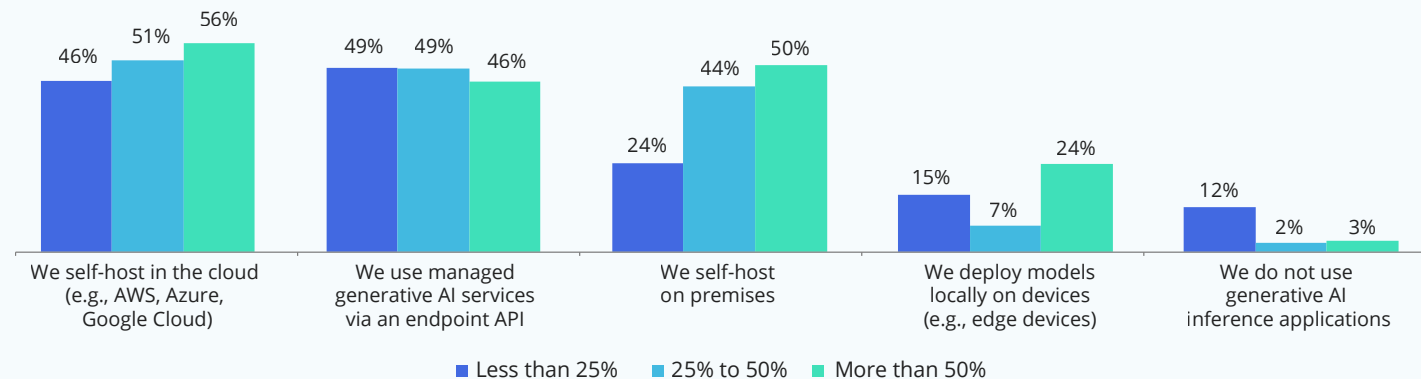
require low latency, such as real-time data processing in IoT environments or in environments with limited connectivity. It can also reduce costs due to the lower reliance on a central model.

Flexible hosting among organizations with higher open source AI usage

Figure 13 shows the hosting preferences for GenAI models for inference among organizations based on the extent of their open source GenAI adoption. It reveals a correlation between open source usage and a preference for flexible, customizable hosting solutions. Specifically, organizations with higher open source adoption (defined as those whose GenAI infrastructure is more than 50% open source) demonstrate a stronger tendency to host on premises and deploy locally on edge devices compared to organizations with lower open source usage.

FIGURE 13: HOSTING LOCATIONS FOR GEN AI MODEL INFERENCE SEGMENTED BY OPEN SOURCE AI USE

Where does your organization host the generative AI models for inference applications? (select all that apply) segmented by **Approximately how much of your organization’s code infrastructure that supports generative AI initiatives is currently from open source?** (select one)



2024 GenAI survey, Q38 by Q19, Sample Size = 170, Valid Cases = 170, Total Mentions = 286

Organizations with greater open source usage are significantly more likely to self-host on premises (50%) compared to those with limited open source adoption (24%). This could imply that heavy open source users likely possess the technical expertise and infrastructure to manage complex deployments independently.

High open source users also show a strong preference for deploying models locally on edge devices with 24% opting for this approach, compared to just 7% among those with minimal open source adoption. Edge deployment is ideal for applications requiring low latency or offline capabilities, which may be more achievable with open source models optimized for edge performance. The data reinforces that open source adoption is closely linked with a preference for flexible, autonomous hosting strategies. A philosophy for a decentralized GenAI deployment might underlie this trend.

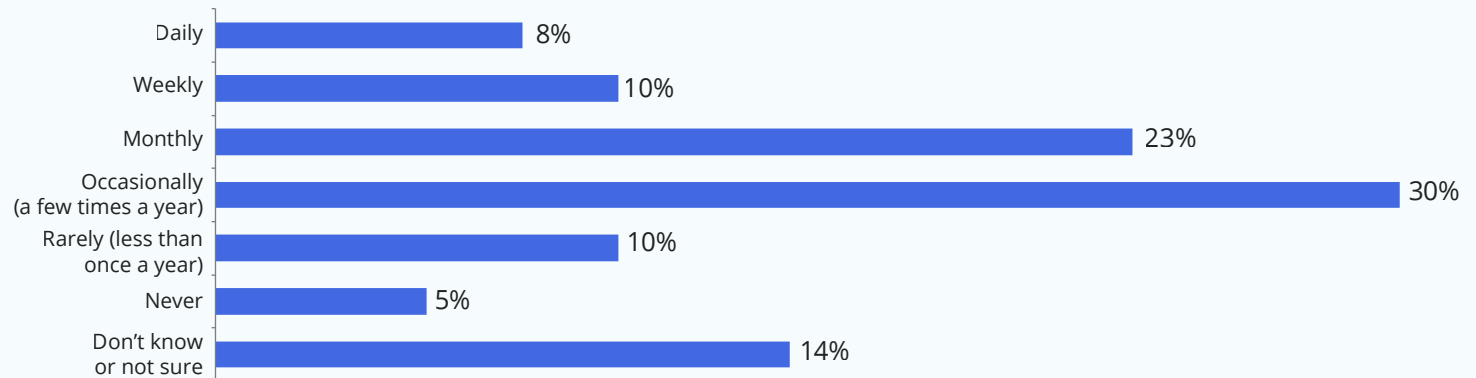
Kubernetes as a key enabler for hosting scalable GenAI workloads

Building on the importance of cloud native and hybrid cloud strategies, Kubernetes emerges as a key enabler for hosting scalable GenAI workloads. It provides the orchestration and flexibility needed to manage complex deployments, especially for organizations with high deployment frequency, as illustrated in this section.

Figure 14 examines the frequency with which organizations deploy GenAI models into production for inference, revealing a diversity in deployment practices and a clear need for scalable infrastructure. While 41% of organizations deploy models frequently (either daily, weekly, or monthly), 30% deploy only occasionally, and 10% rarely deploy at all. This variation in deployment frequency shows the importance of infrastructure that can accommodate both continuous and less frequent model updates.

FIGURE 14: FREQUENCY OF GEN AI DEPLOYMENT FOR INFERENCE

How frequently does your organization deploy generative AI models into production for inference? (select one)



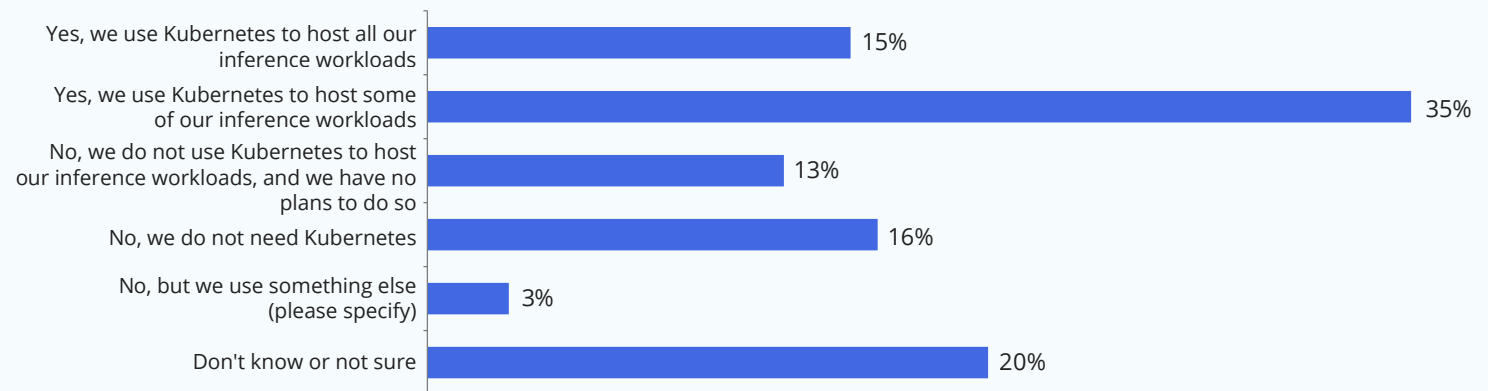
2024 GenAI survey, Q41, Sample Size = 209, answered by organizations who self-host in Q38 or Q39

Organizations that update models regularly demonstrate the operational maturity to manage constant model iterations and optimizations. Frequent updates can be critical in GenAI due to the need for refined outputs and rapid incorporation of new data. Kubernetes provides a robust solution for these organizations by supporting automated scaling, rolling updates, and simplified deployment pipelines. With Kubernetes, organizations can manage these frequent deployments more seamlessly, ensuring models remain up to date without disrupting services.

Figure 15 shows Kubernetes' role in GenAI model hosting, revealing that it has gained strong traction in this space. Half of the respondents (50%) use Kubernetes to host all or some of their GenAI inference workloads. This adoption pattern aligns with Kubernetes' strengths in managing dynamic, resource-intensive applications, which GenAI models typically represent.

FIGURE 15: KUBERNETES USE FOR GEN AI MODEL INFERENCE WORKLOADS

Do you use Kubernetes to host your generative AI model inference workloads? (select all that apply)



2024 GenAI survey, Q40, Sample Size = 208, Valid Cases = 208, Total Mentions = 209, answered by organizations who self-host in Q38 or Q39

Many organizations using Kubernetes only for part of their workloads (35%) may be leveraging a hybrid cloud strategy or using Kubernetes in conjunction with other technologies. This connects to CNCF's support for hybrid and multi-cloud environments, where Kubernetes can manage AI workloads across cloud and on-premises infrastructures seamlessly.

A smaller segment (15%) relies exclusively on Kubernetes for all inference workloads, indicating mature and Kubernetes-

centric infrastructure strategies. These organizations are likely to benefit most from Kubernetes' features such as automated scaling, load balancing, and fault tolerance.

“Organizations value Kubernetes' infrastructure-agnostic capabilities, allowing them to deploy GenAI models across various environments.”

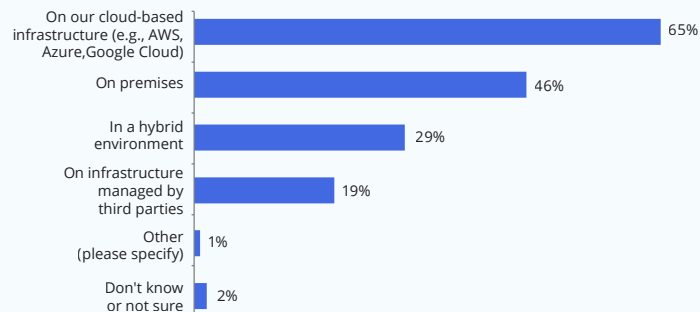
Cloud-based infrastructure leads the way in GenAI model building, with hybrid and on-premises solutions remaining key

With Kubernetes enabling scalable and efficient deployment of GenAI models, the focus shifts to the critical stages of model building and training—and where organizations choose to host these resource-intensive processes.

Figure 16 illustrates that the cloud is the most common infrastructure for GenAI model building, with 65% of organizations relying on cloud platforms. The popularity of cloud solutions is partly due to their elasticity, access to specialized hardware (e.g., GPUs, TPUs), and scalability, which allow organizations to train complex models without the need for heavy upfront infrastructure investments. The cloud's pay-as-you-go model also enables organizations to scale resources according to project demands, which is necessary given the resource-intensive nature of model building and training.

FIGURE 16: HOSTING LOCATION FOR GEN AI MODEL BUILDING AND TRAINING

Where does your organization host generative AI model building and training? (select all that apply)



2024 GenAI survey, Q39, Sample Size = 113, Valid Cases = 113, Total Mentions = 183, answered by organizations who train or build models in Q32

However, nearly half of the respondents (46%) continue to use on-premises infrastructure, often driven by requirements around data security, privacy, or latency-sensitive applications, especially when organizations need full control over data and processing resources. However, the reliance on on-premises infrastructure could also signal that these organizations have less maturity in cloud adoption and so leverage their own hardware.

There is a need for flexible architectures that can balance cost, compliance, and computational needs, and the use of hybrid environments (29%) shows this. A hybrid approach allows sensitive workloads to be kept on premises while offloading less sensitive or resource-intensive tasks to the cloud. This multi-environment approach is important for organizations needing to benefit from cloud scalability while maintaining data control where required. These organizations can leverage Kubernetes and containerized applications to make seamless integration across environments more feasible.

“A hybrid approach allows sensitive workloads to be kept on premises while offloading less sensitive or resource-intensive tasks to the cloud. This multi-environment approach is important for organizations needing to benefit from cloud scalability while maintaining data control where required.”

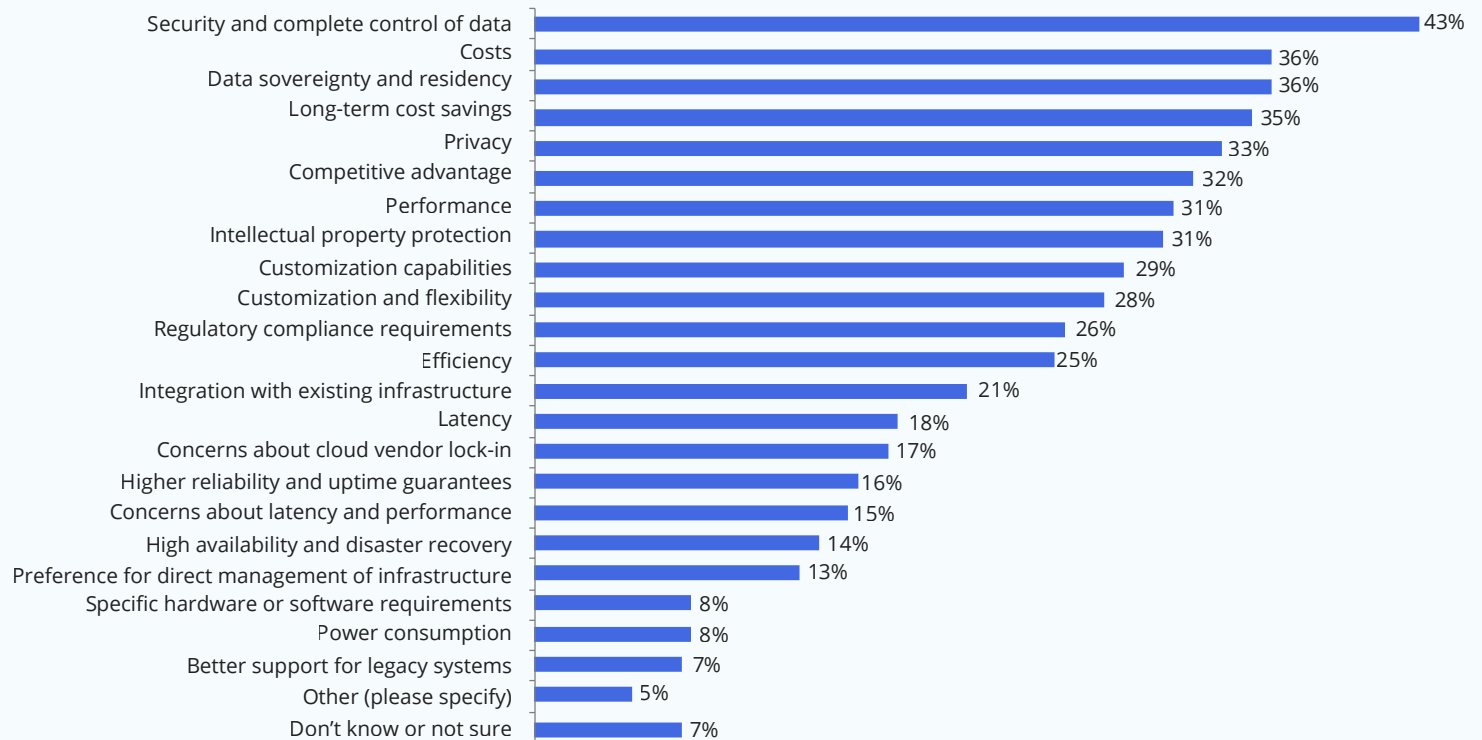
Motivations behind building out GenAI infrastructure

Figure 17 shows why organizations decide to build out GenAI infrastructure in the first place. Top motivations include security and data control (43%), cost (36%), data sovereignty (36%), and

long-term savings (35%). Self-hosting or dedicated infrastructure provides a controlled environment, reducing exposure to breaches and allowing organizations to implement customized security measures.

FIGURE 17: MOTIVATIONS FOR BUILDING OUT GEN AI INFRASTRUCTURE

What are your organization's primary motivations for building out generative AI infrastructure? (select all that apply)



2024 GenAI survey, Q42, Sample Size = 209, Valid Cases = 209, Total Mentions = 1,115, answered by organizations who self-host in Q38 or Q39

Cost considerations are also significant, as organizations balance the high upfront investment of in-house infrastructure against the ongoing costs of public cloud. For some, the predictability of in-house solutions outweighs the potentially rising costs of scalable cloud resources, particularly for compute-

intensive GenAI. Additionally, compliance with data residency laws remains a key driver, with many opting for localized infrastructure to meet strict regional requirements more easily than with a multi-regional public cloud setup.

Challenges in GenAI adoption

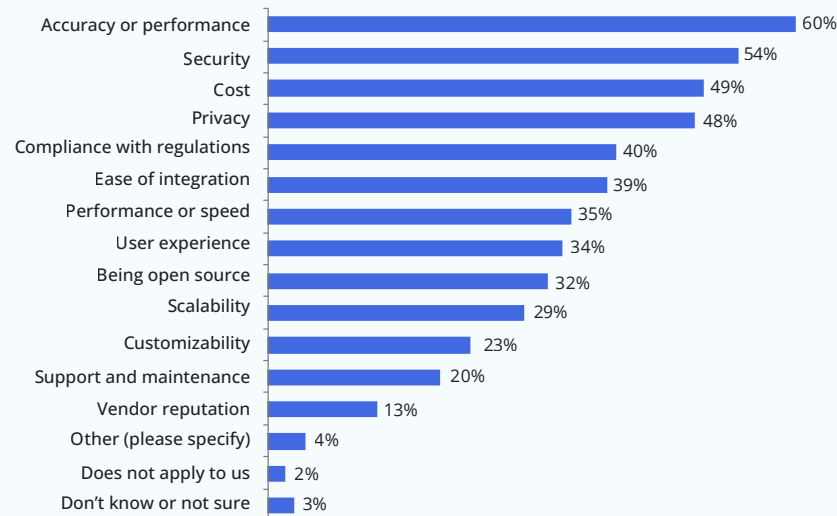
“In my personal opinion, AI is currently massively overhyped by companies. It’s still nowhere near in a state where I would trust it to be a helpful tool for production systems.”

Given these strong motivations for building out dedicated GenAI infrastructure, organizations still face a range of challenges as they implement and scale these solutions. The next section explores these challenges in depth.

Figure 18 shows that organizations prioritize accuracy, performance, security, cost, and privacy when selecting GenAI models or tools. Accuracy and performance rank highest, with 60% of respondents valuing these factors, followed by security (54%), cost (49%), and privacy (48%). These choices reflect a strong need for reliable, secure, and cost-effective solutions that also protect sensitive data.

FIGURE 18: MOST IMPORTANT FACTORS WHEN CHOOSING A GEN AI MODEL OR TOOL

What are the most important characteristics your organization considers when choosing a generative AI model or tool? (select all that apply)



2024 GenAI survey, Q17, Sample Size = 295, Valid Cases = 295, Total Mentions = 1,427, answered by organizations who adopted GenAI in Q7, optional question

Segmented by organization type:

	Cross-industry IT organizations	Industry-specific organizations
1	Accuracy/performance (60%)	Security (61%)
2	Cost (52%)	Privacy (60%)
3	Security (51%)	Accuracy/performance (59%)

2024 GenAI survey, Q17 by Q4, Sample Size = 271, Valid Cases = 271, Total Mentions = 1,326

Priorities vary between sectors: industry-specific organizations focus more on security (61%) and privacy (60%), which regulatory requirements drive in fields such as healthcare and finance. Cross-industry IT organizations, meanwhile, place greater importance on accuracy, performance (60%), and cost (52%) to improve operational efficiency and return on investment. These differing needs reveal varied challenges organizations face in adopting GenAI.

Primary concerns of GenAI adoption

Figure 19 aligns closely with Figure 18, where accuracy, performance, security, and cost emerged as top considerations for choosing generative tools. Here, output quality, privacy, operational costs, and security risks are the key concerns. It is clear that organizations prioritize these factors when selecting tools and also view them as ongoing challenges during the adoption process.

FIGURE 19: PRIMARY CONCERNS OF ADOPTING GEN AI MODELS AND TOOLS

When adopting generative AI models and tools, what are your primary concerns? (select all that apply)



2024 GenAI survey, Q18, Sample Size = 316, Valid Cases = 316, Total Mentions = 1,934

Segmented by organization type:

	Cross-industry IT organizations	Industry-specific organizations
1	Quality of AI output (48%)	Privacy of our data (57%)
2	Cost of operations (42%)	Quality of AI output (53%)
3	Privacy of our data (40%)	Security risks (47%)

2024 GenAI survey, Q18 by Q4, Sample Size = 292, Valid Cases = 292, Total Mentions = 1,774

Concerns differ by industry type. Industry-specific organizations, such as those in finance and healthcare, emphasize privacy (57%) and security (47%) due to strict regulatory and data protection requirements. In contrast, cross-industry IT organizations place the greatest importance on output quality (48%) and operational costs (42%), likely reflecting a focus on delivering dependable and affordable solutions across diverse business environments.

Investment in GenAI

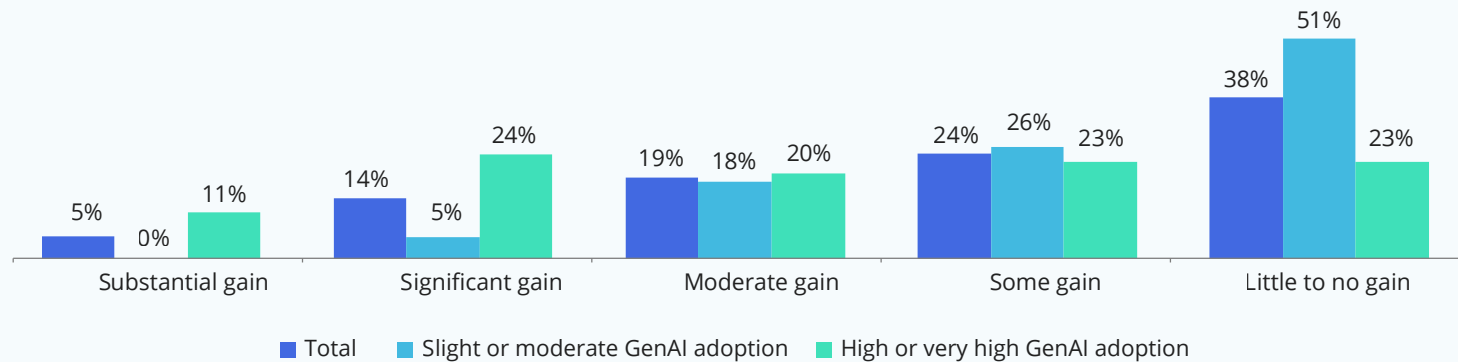
Following the discussion on the key factors and concerns influencing generative tool adoption, Figure 20 examines how these investments translate into revenue gains for organizations. 62% of organizations report seeing some level of revenue gain from their investments in generative tools and models, although

the degree of financial impact varies widely. Organizations with high or very high adoption levels are more likely to report moderate to significant gains, with 24% of this group indicating significant revenue impact. In contrast, those with only slight or moderate adoption levels are more likely to report little to no revenue gain, with 51% in this category.

This data suggests that a stronger commitment to generative tools often correlates with better financial returns. This is likely due to deeper integration into core processes, more refined applications, and a greater focus on measurable outcomes. For organizations still in the early stages of adoption, the results indicate a need to explore strategic integration points where generative tools can add tangible value.

FIGURE 20: REVENUE GAIN FROM GEN AI INVESTMENT SEGMENTED BY GEN AI ADOPTION LEVEL

How much of your organization's investment in generative AI has been converted into revenue gain?
(select one) segmented by **To what extent does your organization adopt generative AI?** (select one)



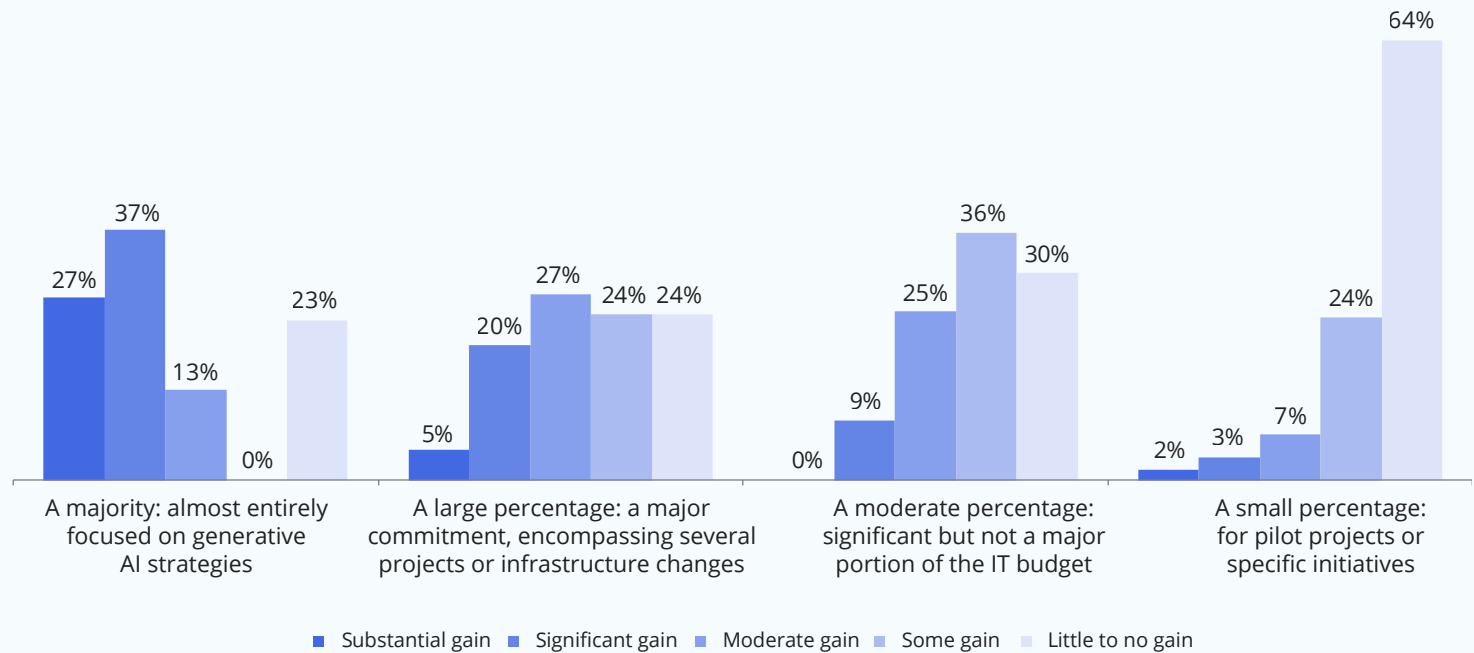
2024 GenAI survey, Q11 by Q7, Sample Size = 234 (DKNS and NA excluded)

Figure 21 shows that while larger investments in generative technology can lead to higher gains, substantial spending does not always guarantee substantial returns. Organizations that

dedicate a majority of their IT budgets to these tools report mixed results: 37% achieve significant gains, but 13% see little to no revenue impact.

FIGURE 21: INVESTMENT IN GEN AI SEGMENTED BY REVENUE GAIN FROM GEN AI

How much is your organization planning to invest in its generative AI strategies during the next 12 months as a percentage of its overall IT budget? (select one) segmented by **How much of your organization's investment in generative AI has been converted into revenue gain?** (select one)



2024 GenAI survey, Q10 by Q11, Sample Size = 224

Smaller investments, such as pilot projects, tend to produce modest gains, with 64% of organizations in this category reporting little to no revenue increase. This suggests that while focused, larger-scale investments have more potential to drive meaningful results, their success depends on how strategically

organizations use these funds. Simply investing more does not ensure high returns; the application, integration, and alignment with business goals are likely critical factors in achieving substantial gains.

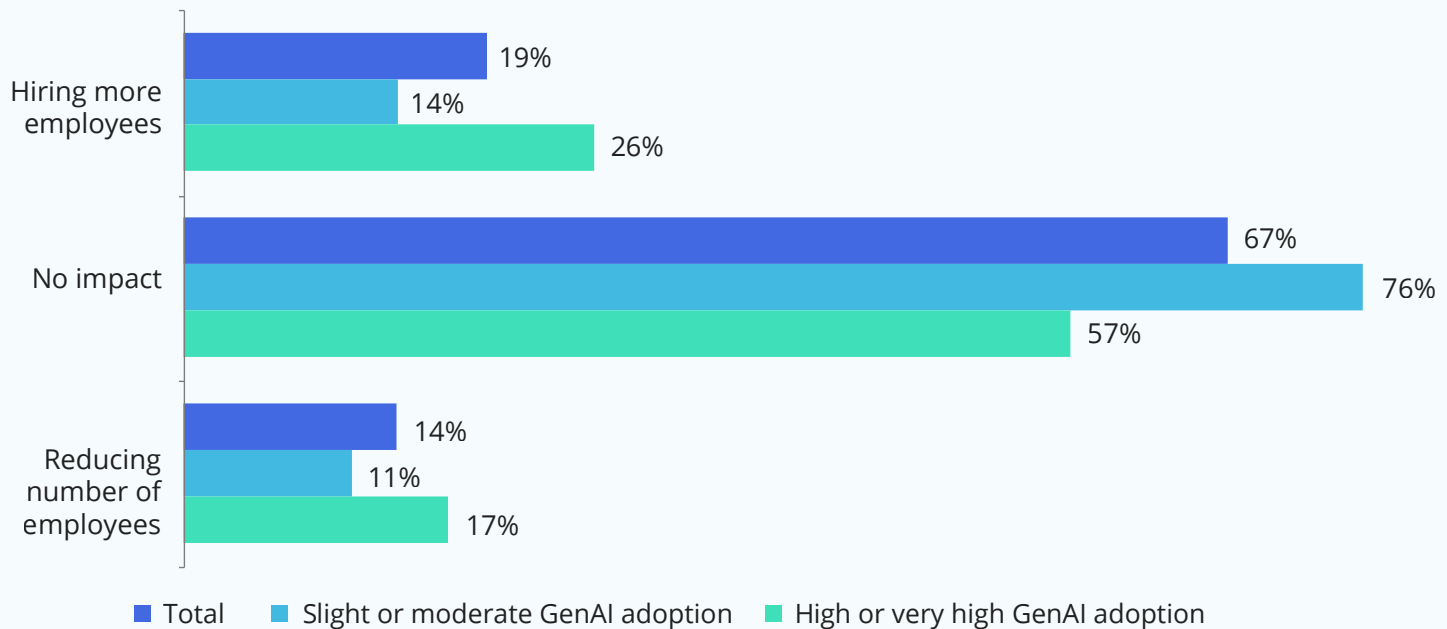
Impact on employment

Figure 22 shows that the majority of organizations (67%) report no impact on employment from adopting generative technologies. This suggests that, for most, these tools are augmenting existing workflows rather than replacing jobs. However, organizations with high levels of adoption show a greater likelihood of both hiring more employees (26%) and, in some cases, reducing staff (17%), indicating that the impact on employment varies with the depth of integration.

For high adopters, increased hiring may reflect the need for specialized skills to support and maintain generative tools or to develop new roles focused on leveraging these technologies effectively. In contrast, staff reductions could indicate efficiency gains in certain areas, where automation allows organizations to streamline specific tasks. Overall, the data points to a balanced impact on employment, with most organizations experiencing little change, while those with more intensive adoption see both opportunities for workforce expansion and potential restructuring.

FIGURE 22: IMPACT OF GEN AI ON EMPLOYMENT COUNT

How has the adoption of generative AI impacted employment in your organization? (select one)
segmented by **To what extent does your organization adopt generative AI?** (select one)



2024 GenAI survey, Q12 by Q7, Sample Size = 257

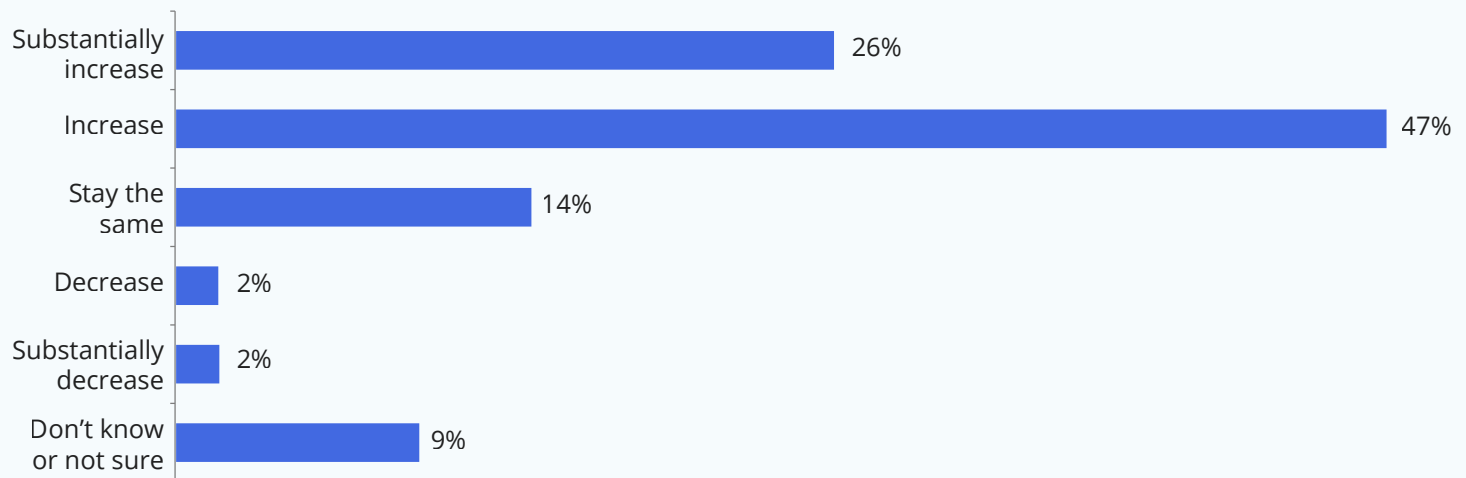
The future of GenAI is open

As we look toward the future of GenAI technology, Figure 23 reveals strong momentum for open source tools in this space. 73% of organizations expect to increase their use of open source generative tools over the next two years, with 26% anticipating a substantial rise. In the next section, we will explore the priorities organizations have for open source projects, along with key perspectives and shared goals for the future of GenAI.

“AI is going to be the next future platform; learn with it, and live it.”

FIGURE 23: INCREASE OF OPEN SOURCE GEN AI USE IN THE NEXT TWO YEARS

How do you expect this use of open source generative AI in your organization to change in the next two years? (select one)



2024 GenAI survey, Q20, Sample Size = 297, answered by organizations who adopted GenAI in Q7

Top priorities for open source projects

As shown in Figure 24, the top three priorities for open source projects over the next year are security, open standards, and data governance, each reflecting core needs as organizations increasingly rely on these tools.

Security is a high priority, as generative tools often process sensitive data, requiring strong protections against unauthorized access and vulnerabilities. CNCF's KubeFlow, for example, is designed to make model deployment and management easier on Kubernetes, integrating security practices throughout the deployment process.

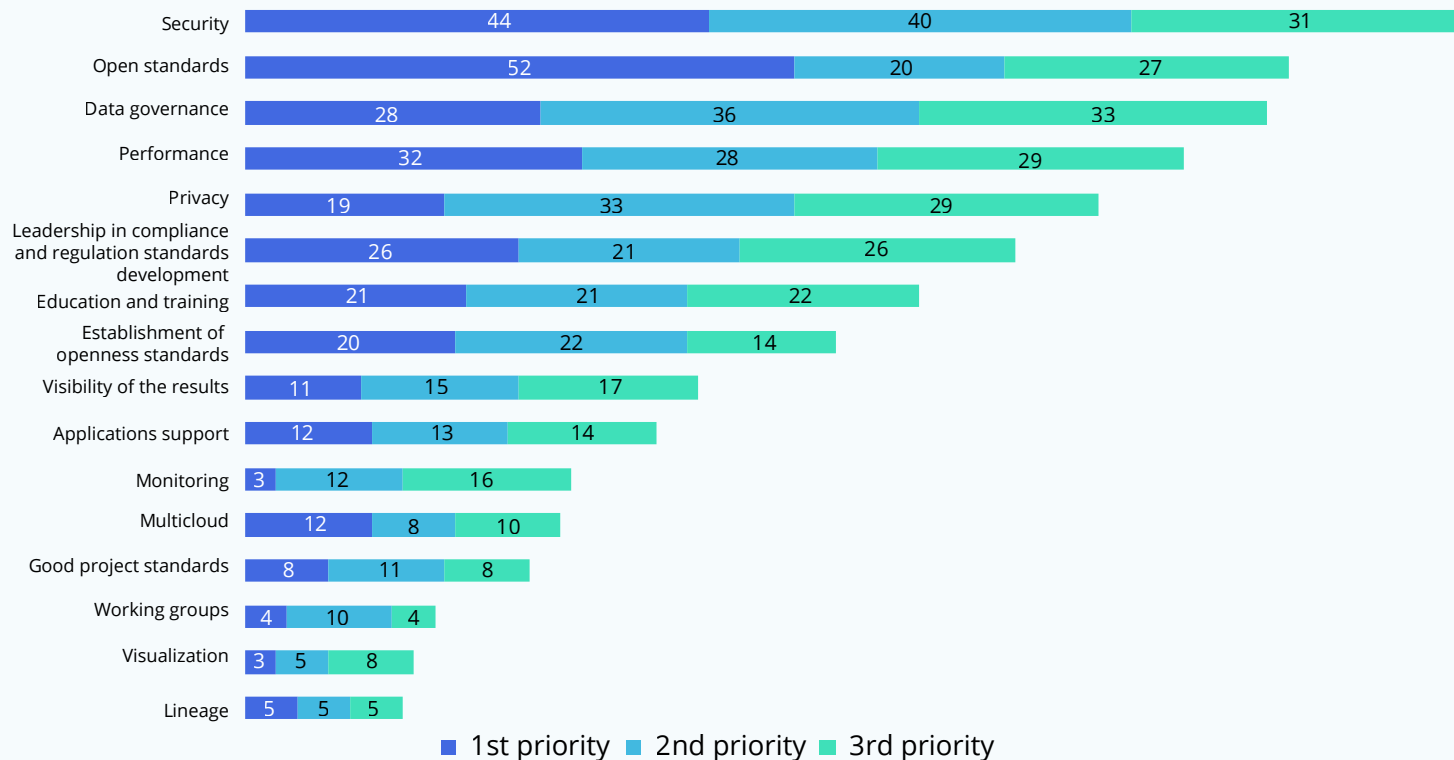
Open standards enable interoperability, allowing organizations to easily integrate generative tools with other systems, whether for data processing, model deployment, or visualization. For example, Open Neural Network Exchange establishes open standards for model interoperability, which means it allows models trained in one framework to be run in another.

and compliance. For example, ODPi Egeria provides an open framework for managing metadata and establishing clear data lineage, which is vital when working with the large, complex datasets typical in GenAI technology. Delta Lake also supports data governance as a common architecture for storing the vast datasets needed for generative processes.

Data governance ensures that organizations manage data used in generative workflows responsibly, with attention to quality

FIGURE 24: PRIORITIES FOR OPEN SOURCE AI PROJECTS

What would be your top three priorities for open source AI projects to focus on in the next year?



2024 GenAI survey, Q30, Sample size = 15-115, respondents identified up to three priorities / rows, optional question

The role of open source AI in the future of GenAI

Open source tools and frameworks are not only foundational today but are essential for the future of GenAI. According to Figure 25, 82% of respondents agree that open source solutions are crucial for a positive future, with 61% believing the benefits outweigh the risks.

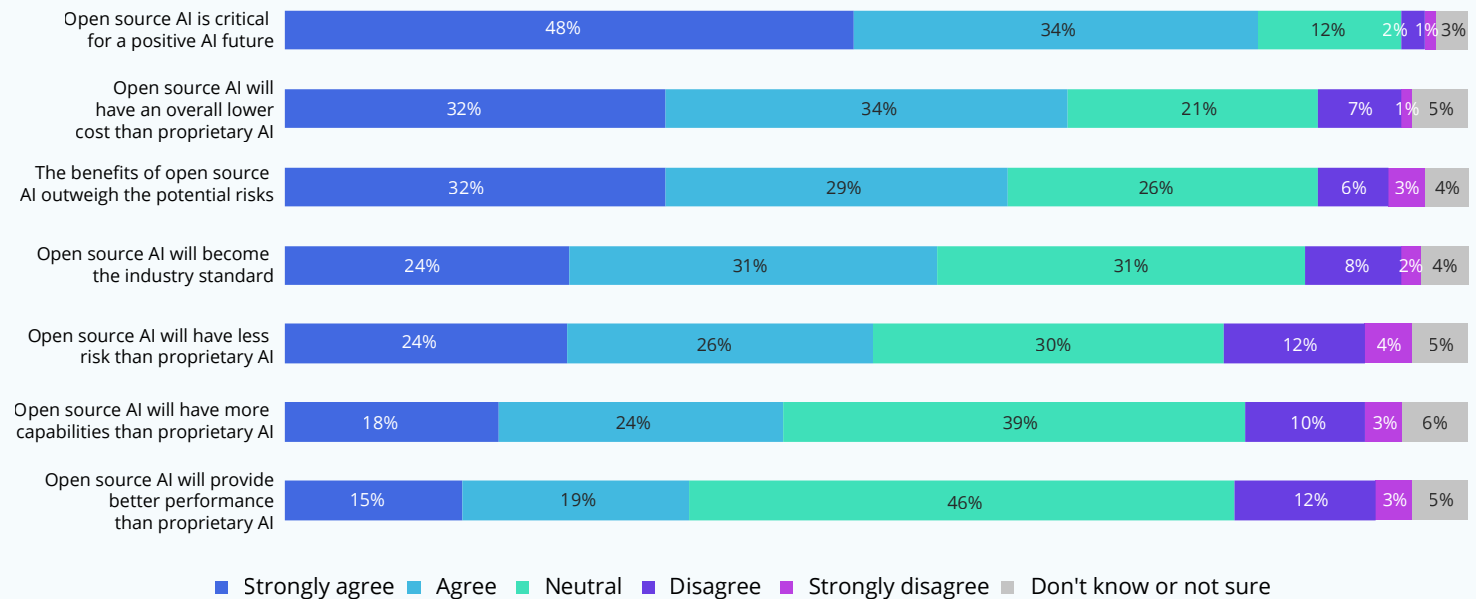
Cost is another key factor driving the open source approach. 66% percent of respondents believe that open source tools

will provide a lower overall cost compared to proprietary options, making this path attractive for both large and small organizations. Community-driven projects reduce licensing fees and allow for continual improvement, giving companies access to high-quality solutions without heavy financial commitments.

Looking ahead, 48% of respondents expect open source solutions to become the industry standard. This reflects a strong demand for critical tools, frameworks, and infrastructure to be accessible to all, allowing companies to build, deploy, and scale solutions without being locked into proprietary ecosystems.

FIGURE 25: PERCEPTIONS ON THE BENEFITS OF OPEN SOURCE AI

To what extent do you agree with the following statements? (one response per row)



2024 GenAI survey, Q28, Sample size = 311-314, not all rows required

Building on the strong belief that open source is foundational to GenAI, Figure 26 further verifies the demand for openness in AI. Figure 26 shows strong agreement among respondents (83%) that AI needs to become increasingly open.

“Openness will help organizations and developers build trust and adapt solutions more effectively to diverse applications and needs.”

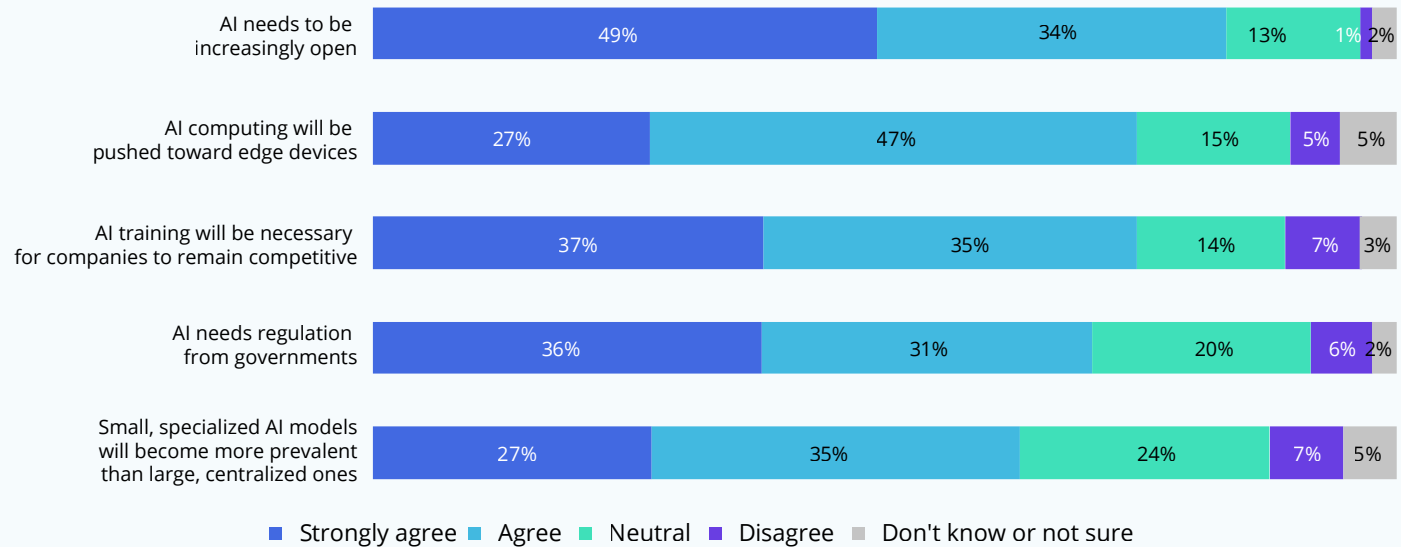
Considering aspects other than openness, 74% of respondents forecast a push toward edge devices, which reflects the growing

importance of decentralized processing. Running models closer to where data is generated allows for faster response times, reduces dependency on central servers, and helps privacy by keeping data local.

Over two-thirds (67%) of respondents feel that government regulation is necessary, a sign that the industry sees a role for oversight in ensuring responsible development and use. Finally, 62% predict a shift toward smaller, more specialized models rather than large, centralized ones. Altogether, these views suggest a future where generative technology is not only more open and transparent but also increasingly distributed, customizable, and responsive to both regulatory and operational demands.

FIGURE 26: PERCEPTIONS ON THE FUTURE OF GEN AI

Regarding the future of generative AI, to what extent do you agree with the following statements? (one response per row)



2024 GenAI survey, Q45, Sample size = 312-316, not all rows required

Conclusions and recommendations

“It’s still a watch this space. I believe GenAI does have significant potential, but the impression I get is that it is currently at the peak of the hype cycle.”

The findings of this report highlight the vital role that open source technology plays in the rapid adoption and democratization of GenAI. Open source frameworks, tools, and communities are making advanced GenAI models accessible to organizations, enabling a foundation for tailored solutions and fostering a collaborative ecosystem. High adoption rates among organizations with open source foundations emphasize the adaptability, cost-effectiveness, and transparency that open source provides.

Organizations with a high reliance on open source GenAI report enhanced productivity, scalability, and innovation. Support from communities such as the Linux Foundation and CNCF has been instrumental in advancing open source GenAI, providing shared standards, robust frameworks, and secure infrastructures. Open source tools, such as TensorFlow, PyTorch, LangChain, and LlamaIndex, are foundational in enabling organizations to develop, deploy, and manage GenAI models effectively. These tools not only accelerate the adoption of GenAI but also mitigate costs, reduce dependencies on proprietary systems, and promote a community-driven model of continuous improvement.

As GenAI technologies mature, open source will likely continue to be essential for their responsible and efficient use. Open source infrastructure, combined with cloud native architectures,

enables scalable and flexible deployments that align with organizational goals. Kubernetes, in particular, has emerged as a critical tool for orchestrating and managing GenAI workloads, providing the required scalability, flexibility, and operational resilience. The integration of open source solutions allows organizations to maintain control over data and model behaviors, supporting regulatory compliance and ethical considerations.

Moving forward, it is recommended that organizations prioritize open source solutions as a core strategy for GenAI adoption. Moreover, active contribution to open source projects should be encouraged, as this engagement not only strengthens the broader AI community but also allows organizations to shape tools and frameworks to better meet their specific needs.

Looking to the future of GenAI technology, we expect open source’s influence in the AI domain to expand further. This survey reveals that 83% of organizations strongly agree or agree that AI needs to be increasingly open. Additionally, 82% report that open source AI is a critical component for a sustainable AI future, with 61% expressing confidence that the benefits of open source outweigh the associated risks. Growth of open source GenAI technology is likely to be significant with 73% of organizations expecting to increase their use of open source generative tools over the next two years, with 26% anticipating a substantial rise in use.

This survey reveals that 82% of respondents view open source as a critical component for a sustainable AI future, with 61% expressing confidence that the benefits of open source outweigh the associated risks. By fostering an environment of shared innovation, transparency, and cost efficiency, open source

initiatives are likely to become the industry standard for AI development. In particular, community-driven improvements in areas such as model security, ethical governance, and interoperability will ensure that organizations deploy GenAI technologies responsibly and inclusively across sectors.

To address the complexities and ethical considerations surrounding GenAI, establishing robust open governance practices within open source communities is essential. This includes adopting transparent standards for data handling, model training, and operational practices, ensuring the AI systems align with broader societal values. Additionally, open source-based models and frameworks should be supported by

neutral organizations, which can provide impartial governance and foster inclusive decision-making, especially as GenAI applications continue to evolve.

In summary, open source will remain indispensable for the future of GenAI, driving accessibility, innovation, and ethical considerations. Organizations that embrace open source strategies and contribute to these communities are well-positioned to remain at the forefront of AI advancements. By fostering an ecosystem that prioritizes shared knowledge, responsible governance, and continuous innovation, the industry can ensure that GenAI technology develops in ways that benefit organizations, individuals, and society as a whole.



Methodology

About the survey

A web survey that Linux Foundation Research and its partners conducted from August to September 2024 provided the basis for this study. The survey's goal was to understand the deployment, use, and challenges of GenAI technologies in organizations and the role of open source in this domain.

We sourced our sample from Linux Foundation subscribers, members, partner communities, and social media. We addressed data quality through extensive prescreening, survey screening questions, and data quality checks to ensure that respondents had sufficient professional experience to answer questions accurately on behalf of the organization they worked for.

We collected survey data from industry-specific companies, IT vendors and service providers, and nonprofit, academic, and government organizations. Respondents spanned many vertical industries and companies of all sizes, and we collected data from four geographies, including the Americas, Europe, Asia-Pacific, and the Rest of World.

The 2024 GenAI survey comprised 46 questions that addressed screening, respondent demographics, GenAI investment, GenAI concerns, open source AI in organizations, GenAI techniques, GenAI infrastructure, and gains from GenAI. For information about access to the 2024 GenAI survey, its dataset, and survey frequencies, see the Data.World access information below. The high-level design of the survey is shown in Figure 27.



FIGURE 27: SURVEY DESIGN

Questions	Question categories	Who answers the questions
Q1 - Q6	Demographics	All respondents (N=316)
Q7 - Q16	Importance and investment in GenAI in your org	Orgs who adopted GenAI (N=297)
Q17 - Q18	GenAI concerns	Orgs who adopted GenAI (N=297 for Q17, All respondents (N=316) for Q18)
Q19 - Q26	Open source AI in your org	Orgs who adopted GenAI (N=297)
Q27 - Q30	Perspectives on open source AI	All respondents (N=316)
Q31 - Q34	GenAI techniques	Orgs who adopted GenAI (N=297)
Q35 - Q39	Infrastructure for GenAI	Orgs who use GenAI for inference (N=193) Q35, Q36, Q38, or who train models (N=113), Q37, Q39
Q40 - Q42	Self-hosting GenAI	Orgs who self-host (N=209)
Q43 - Q44	Individual gains from GenAI	Orgs who adopted GenAI (N=297)
Q45 - Q46	Future of GenAI	All respondents (N=316)

Survey respondent screening involved the use of three variables to validate the respondent:

- Must be familiar, very familiar, or extremely familiar with the adoption of GenAI intelligence in your organization
- Must work for an organization
- Must have some level of professional experience

A total of 316 respondents completed the survey. The margin of error for this sample size was +/- 4.7% at a 90% confidence level and +/- 5.5% at a 95% confidence level.

The percentage values in this report may not total to exactly 100% due to rounding.

Although there was a requirement for respondents to answer nearly all questions in the survey, we made a provision when a respondent was unable to answer a question by adding a “Don’t know or not sure” (DKNS) response to the list of responses for every question. However, this created a variety of analytical challenges. One approach was to treat a DKNS just like any other response to know the percentage of respondents that answered DKNS. The advantage of this approach is that it shows the exact distribution of data collected. The challenge with this approach is that it can distort the distribution of valid responses, i.e., responses where respondents could answer the question. Some of the analyses in this report exclude DKNS responses. This is because we can classify the missing data as either missing at random or missing completely at random. Excluding DKNS data from a question does not change the distribution of data (counts) for the other responses, but it does change the size of the denominator used to calculate the percentage of responses across the remaining responses. This has the effect of proportionally increasing the percentage values of the remaining responses. Where we have elected to exclude DKNS data, the footnote for the figure includes the phrase, “DKNS responses excluded from the analysis.”

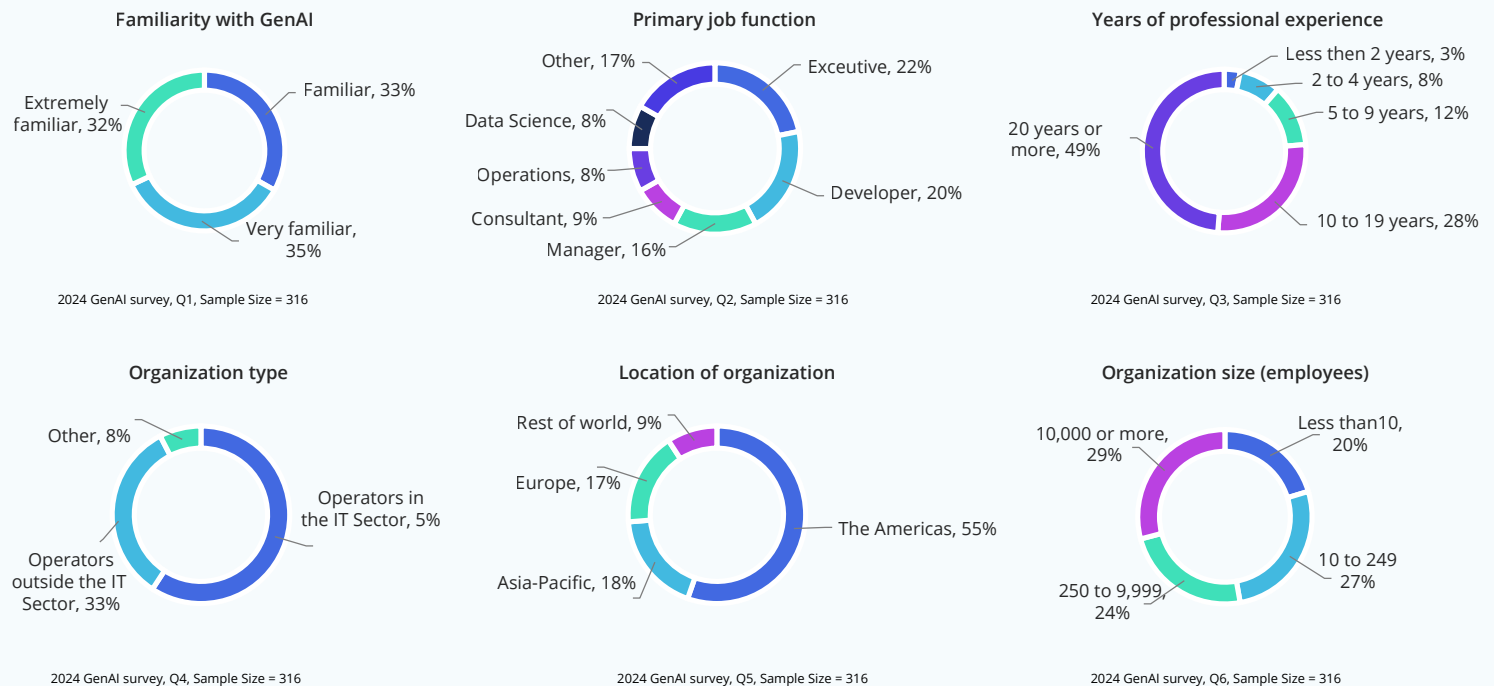
Data.World access

Linux Foundation Research makes each of its empirical project datasets available on Data.World. This dataset includes the survey instrument, raw survey data, screening and filtering criteria, and frequency charts for each question in the survey. You can find Linux Foundation Research datasets, including this project, at data.world/thelinuxfoundation. Access to Linux Foundation datasets is free but does require you to create a Data.World account.

Respondent demographics

These demographics provide you with a profile of the 2024 GenAI survey respondents. We have regrouped all the demographics in Figure 28 to facilitate a more insightful analysis. For the original survey instrument, source data, and study frequencies, please see the Data.World access described above.

FIGURE 28: RESPONDENT DEMOGRAPHICS



About the authors

ADRIENN LAWSON is a data analyst at the Linux Foundation. Adrienn obtained a master's degree in social data science from the University of Oxford. She is responsible for survey development, analysis, and report writing. Adrienn has previously conducted research at the University of Oxford, the Budapest Institute for Policy Analysis, and the U.K.'s Office for National Statistics. She has a strong fascination with the collective power of open source collaboration within geographically dispersed communities. Additionally, she is most interested in researching trends and solutions for challenges related to open source software funding, sustainability, and supporting developers in their pursuit of responsible technological advancement.

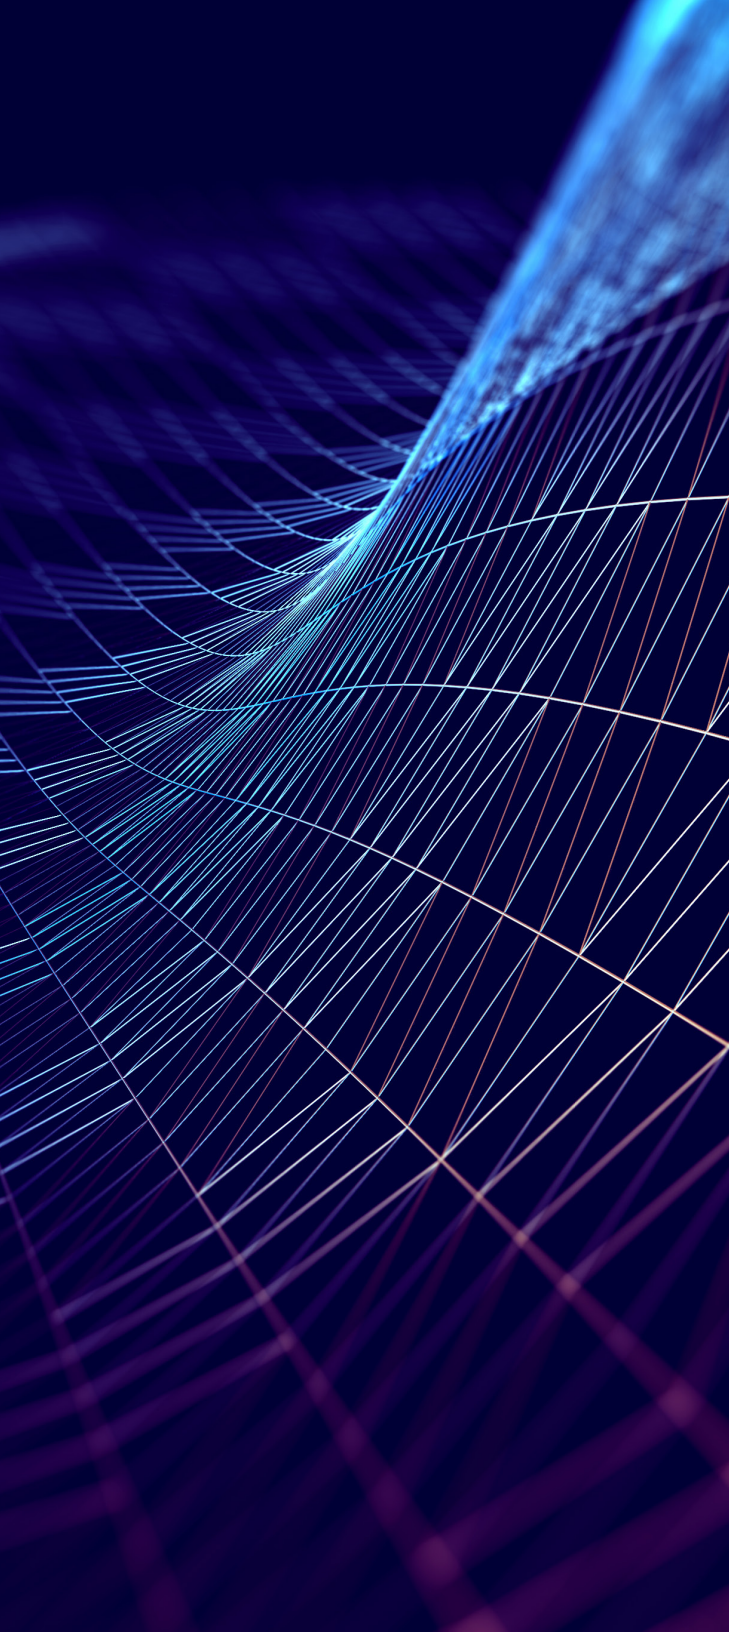
STEPHEN HENDRICK is the vice president of research at the Linux Foundation, where he is the principal investigator on a variety of research projects core to the Linux Foundation's understanding of how open source software is an engine of innovation for producers and consumers of IT. Steve specializes in primary research techniques developed over 30 years as a software industry analyst. Steve is a subject-matter expert in application development and deployment topics, including DevOps, application management, and decision analytics. Steve brings experience in a variety of quantitative and qualitative research techniques that enable deep insight into market dynamics and has pioneered research across many application development and deployment domains. Steve has authored over 1,000 publications and provided market guidance through syndicated research and custom consulting to the world's leading software vendors and high-profile start-ups.

NANCY RAUSCH is a senior program manager for data science at the Linux Foundation, where she specializes in open source

AI technologies. She has extensive experience in the practical application of AI in various industry sectors, including data quality, data governance, renewable energy forecasting, and healthcare analytics. She sits on the Advisory Board for the NSF-funded non-profit LASER Institute, which promotes Learning Analytics in STEM Education Research. She has authored many research papers and publications in AI and data science. She holds two Master of Science degrees: one in business analytics from Capella University and another in statistics and computer engineering from Duke University. Additionally, she earned her Bachelor of Science in electrical engineering from Michigan Technological University.

JEFFREY SICA is head of projects at the CNCF, with a focus on improving maintainer experience, building communities, and project automation. Before that, he worked at Red Hat and the University of Michigan focusing on cloud native technologies and CICD patterns. Jeffrey has been a contributor to upstream Kubernetes, helping in SIG-Contribex, SIG-Release, and SIG-UI. He passionately advocates for open source development and recognizing and alleviating burnout.

MARCO GEROSA is a full professor of computer science at Northern Arizona University and a research analyst at LF Research. His research on software engineering and open source software has resulted in over 300 publications in top-tier venues. He serves on the program committee of renowned conferences and as a reviewer for several journals. Dr. Gerosa has a Ph.D., a master's in informatics, and a B.S. in computer engineering. He is a senior member of the Institute of Electrical and Electronics Engineers and the Association for Computing Machinery. He supervised several Ph.D. and M.Sc. students who are now researchers at top institutions. He also has more than 20 years of teaching experience. For more information, visit <http://www.marcoagerosa.com>.



Acknowledgments

We thank all the participants of the survey for kindly sharing their insights and experience in search and analytics. Special thanks to the peer reviewers and Linux Foundation colleagues for their involvement in the various stages of the research process, including:

- Carly Akerly
- Chris Aniszczyk
- Elizabeth Bushard
- Hilary Carter
- Mia Chaszeyka
- Katie Greenley
- Ibrahim Haddad
- Anna Hermansen
- Christina Oliviero
- Jen Shelby
- Matt White

LF AI & DATA


The LF AI & Data is a global not for profit foundation that hosts critical components of the global AI & Data technology infrastructure. It brings together the world's top developers, end users, and vendors to identify and contribute to the projects and initiatives that address industry challenges for the benefit of all participants.




Cloud native computing leverages an open source software stack to deploy applications as microservices, where each component is packaged into its own container and orchestrated dynamically to optimize resource utilization. The Cloud Native Computing Foundation (CNCF) hosts key projects within the cloud native ecosystem, including Kubernetes, Envoy, Prometheus, and many others. CNCF serves as a neutral hub for collaboration, bringing together leading developers, end users, and vendors—from the world's largest public cloud providers and enterprise software companies to innovative startups. As part of The Linux Foundation, a nonprofit organization, CNCF fosters the growth and adoption of cloud native technologies across industries. For more information, visit www.cncf.io.

 x.com/cloudnativefdn

 youtube.com/c/cloudnativefdn

 facebook.com/CloudNativeComputingFoundation

 github.com/cncf

 linkedin.com/cloud-native-computing




About Linux Foundation Research

Founded in 2021, **Linux Foundation Research** explores the growing scale of open source collaboration, providing insight into emerging technology trends, best practices, and the global impact of open source projects. Through leveraging project databases and networks, and a commitment to best practices in quantitative and qualitative methodologies, Linux Foundation Research is creating the go-to library for open source insights for the benefit of organizations the world over.

 x.com/linuxfoundation

 facebook.com/TheLinuxFoundation

 linkedin.com/company/TheLinuxFoundation



Copyright © 2024 **The Linux Foundation**

This report is licensed under the **Creative Commons Attribution-NoDerivatives 4.0 International Public License**.

To reference this work, please cite as follows: Adrienn Lawson, Stephen Hendrick, Nancy Rausch, Jeffrey Sica, Marco Gerosa, "Shaping the Future of Generative AI: The Impact of Open Source Innovation," foreword by Hilary Carter, The Linux Foundation, November 2024.