



April 2025

Open source technology in the age of AI

With more organizations deploying generative AI across business functions, a new survey finds that leaders are increasingly turning to open source solutions to build out their tech stacks.



This report is part of a research collaboration among McKinsey, the Mozilla Foundation, and the Patrick J. McGovern Foundation.

At a glance

Open source AI usage and trends

- *The use of open source AI technologies is widespread.* More than 50 percent of respondents reported leveraging open source in each of the data, models, and tools areas of the tech stack.
- *Technical maturity and developer experience influence open source use.* Open source AI use is highest in the technology, media, and telecommunications sectors (70 percent), and experienced AI developers are 40 percent more likely to use open source.
- *Organizations are using open source tools from familiar players.* The most commonly used open source tools among enterprises, as of January 2025, are those developed by large technology players, such as Meta's Llama and Google's Gemma.

Value to organizations and developers

- *Most respondents are satisfied with their open source AI models.* About ten times more respondents reported being satisfied than being dissatisfied with their use of open source technology. The top reasons for satisfaction are performance and ease of use.
- *Open source leads on cost benefits, while proprietary tools have faster time to value.* Respondents say that open source AI has lower implementation costs (60 percent of respondents) and lower maintenance costs (46 percent). But respondents see faster time to value from proprietary tools (48 percent).
- *Developers value open source tools.* Most developers (81 percent) report that experience with open source tools is highly valued in their field and that working with such tools is important to their job satisfaction (66 percent).

Future outlook on open source AI

- *Open source use is likely to grow.* Three-quarters of respondents expect to increase their use of open source AI technologies over the next few years.
- *Organizations are open to a mixture of open and proprietary solutions.* Nearly three-quarters of respondents (more than 70 percent) say they are open to either open source or proprietary technologies across areas of the tech stack.

Risks and mitigation strategies

- *Open source AI tools involve potential challenges.* Respondents cite concerns about cybersecurity (62 percent), regulatory compliance (54 percent), and intellectual property (50 percent) when engaging with AI tools.
- *Organizations are implementing safeguards to manage open source AI risks.* Strategies include strengthening information security frameworks and software supply chain controls, third-party evaluation of models, and guardrails to limit model behavior.

Open source software has long been a critical part of the technology ecosystem. Unlike most commercial software, which typically requires a commercial license or subscription and restricts access to its core technology, open source tools are developed collaboratively and made available to the public to use, modify, and distribute with far fewer restrictions, giving developers the ability to adapt and shape well-tailored solutions to the particular needs of their organizations.

The current age of artificial intelligence is no different. As more enterprises build and deploy AI-driven solutions across their businesses, they are turning to a growing array of open source technologies—offerings such as Meta's Llama family, Google's Gemma family, the Allen Institute for Artificial Intelligence's OLMo family, and more recently Nvidia's NeMo family, DeepSeek-R1, and Alibaba's Qwen 2.5-Max—many of which are fast closing the performance gap relative to proprietary AI models.

A new, first-of-its-kind survey of more than 700 technology leaders and senior developers across 41 countries by McKinsey, the Mozilla Foundation, and the Patrick J. McGovern Foundation provides the largest and most detailed analysis of how enterprises are thinking about and deploying open source AI in their organizations. The results suggest that leaders are embracing open source tools as an essential component of their technology stacks, citing advantages including high performance, ease of use, and lower implementation and maintenance costs relative to proprietary tools. Developers, meanwhile, increasingly view experience with open source AI as an important part of their overall job satisfaction. While open solutions come with concerns about data security and time to value, more than three-quarters of survey respondents say they expect to increase their use of open source AI in the years ahead.

In other words, open source should be a key ingredient in every enterprise's AI technology solution strategy. As AI makes its way into nearly every business function, business leaders and technologists alike should embrace the potential of open technologies or risk surrendering a potential source of competitive advantage.

Vilas S. Dhar, president, Patrick J. McGovern Foundation:

“By democratizing access to innovation ecosystems, open source puts the tools of creation into everyone’s hands, not just those with the deepest pockets. This transforms users into builders and consumers into creators. We see this clearly now: Humanity’s most urgent challenges demand collaborative intelligence that crosses borders and disciplines. Innovation accelerates when expertise flows freely, when ideas collide and combine without permission. Forward-thinking means recognizing that our collective imagination will always outpace what any single lab can produce. The future of AI belongs to ecosystems, not empires.”

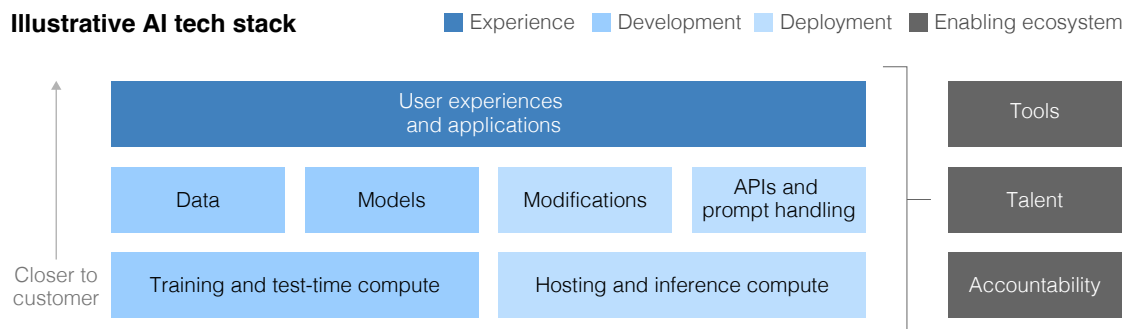
How open source is used in the AI technology stack

The survey included more than 700 respondents across 41 countries, with a particular focus on Brazil, France, India, the United Kingdom, and the United States. Respondents were categorized as either technology decision-makers or senior developers and all have experience working with AI technology systems. (See appendix for respondent details.) Using an AI technology stack adapted from previous Mozilla publications¹ (Exhibit 1), survey respondents were asked about their experiences with open source and proprietary AI technology across seven areas of the AI tech stack: data; models; hosting/inference; modifications; APIs and prompt handling; tools; and user experience/applications. (See definitions and examples of open and “partially open” technologies in the table.) Respondents selected from an extensive list of examples in the survey:

- *Data*: These are data that can be used to train or adapt AI models; includes pretraining, evaluation, fine-tuning, and preference data (such as Mozilla’s Common Voice, the Pile from EleutherAI, and the Allen Institute’s Dolma).
- *Models*: These include the model weights (for example, pretraining, checkpoint, adaptation) as well as the code used to train the model (for example, Mistral, Gemma, Llama, EleutherAI’s GPT-J, Stable Diffusion, and Aya by Cohere).
- *Hosting/inference compute*: This involves servers or cloud infrastructure that host AI models, so that users can run inference (such as Llamafire, NomicAI, Ollama, llama.cpp).
- *Modifications*: This involves fine-tuning, adapters, integrators, and the like, which help adapt foundation models to specific use cases (such as PEFT [parameter-efficient fine-tuning] and LoRa [low-rank adaptation]).
- *APIs and prompt handling*: APIs and handling of prompts make AI usable in specific contexts and domains (for example, Hugging Face Serverless Inference API).
- *Tools*: These are tools that augment the AI development and deployment process, including AI orchestration, security, observability, and evaluation (such as PyTorch, Tensorflow, LangChain, and Llama Guard).
- *User experiences/applications*: These are end-user experiences enabled by AI (such as HuggingChat).

Exhibit 1

An AI tech stack can be built with both open source and proprietary components.



Source: *Public AI: Making AI work for everyone, by everyone*, Mozilla, Sept 30, 2024

¹ *Public AI: Making AI work for everyone, by everyone*, Mozilla, September 2024.

For nonmodel elements of this technology stack (software tools, packages, and data sets), technologies are considered open if they fit standard definitions of open source software, such as being free to use, study, and modify. However, the complexity of AI models, and large language models (LLMs) in particular, has raised questions about the criteria for open source AI models. That said, there is much debate about the definitions and criteria for “openness” in AI. Mozilla supports the definition of open source AI put forward by the Open Source Initiative (OSI),² which includes requirements for openness across data information, code, and parameters. Technologies that meet the OSI standard for “open source” or “open source AI” were considered open source in the survey. Popular tools that have open components but do not meet the OSI standard (for example, open weights models or software with non-OSI-approved open source licenses) were considered partially open in the survey. In charts where the distinction is not made, we include both partially open and open source tools in our open source bucket (table).

How organizations are integrating open source into their technology stacks

Our survey finds that use of open source AI is common across respondents, particularly in areas of the technology stack where it can be seamlessly integrated with existing enterprise systems and security protocols. More than 50 percent of respondents report using open source AI in the data, models, and tools areas of the tech stack (Exhibit 2). Open source is least common in modifications (such as fine-tuning and adapters) and hosting/inference compute; this may be because open source inference projects are still relatively new (for example, one of the most notable new tools, vLLM, was developed at Berkeley’s Sky Computing Lab in April 2024³). It may also indicate that users are choosing to modify their models with internal packages and proprietary data for custom use cases.

Table

AI models fall on a spectrum, from totally closed to completely open.


	Description	Example
Closed	Level 1: Internal/tightly controlled; limited or no external access	Gemini Ultra (initially)
	Level 2: External access (eg, via API); lacking full weights and training information (weights are not open source)	GPT-4 (API), Claude
Partially open	Level 3: Open weights models; code and data limited or not publicly available	Llama 3 and 4
	Level 4: Open weights models with supporting material (eg, code); reasons for not meeting the definition of open source AI established by the Open Source Initiative (OSI) could include having non-OSI-approved licenses or insufficient data	Command R+ (Cohere), Stable Diffusion 2
Open	Level 5: Meets the OSI definition for open source AI (such as having source code, no restrictions on redistribution, and a license that permits modifications and is technology-neutral)	Bloom (BigScience)
	Level 6: Exceeds the OSI standard; data is publicly available	Pythia (EleutherAI), T5 (Google), Dolly 2.0 (Databricks)

² Mozilla Blog, “Celebrating an important step forward for open source AI,” blog entry by Ayah Bdeir, Imo Udom, and Nik Marda, August 22, 2024.

³ Ivan Ortega, “A high-throughput and memory-efficient inference and serving engine for LLMs,” UC Berkeley Sky Computing Lab, April 25, 2024.

Use of open source AI solutions is substantial across the tech stack, with the highest share of respondents reporting it in models and tools.

Organizations' regular use of open source AI solutions, by tech stack area, % of respondents



Source: McKinsey Open Source AI Survey, 703 participants with experience in working with AI tech systems, Dec 9, 2024–Jan 24, 2025

While there are many varieties of open source AI technologies available across the tech stack, organizations are most frequently using those created by major corporations such as Google, Meta, and Microsoft. In models, 61 percent of respondents who reported using open source models had used Meta's Llama, 40 percent had used Google's Gemma, 32 percent had used Mistral, and 28 percent had used Microsoft's Phi. Of the respondents using open source AI tools, the most popular were PyTorch (58 percent), TensorFlow (57 percent), PostgreSQL (45 percent), and LangChain (33 percent). With the shift to agent development, libraries that provide open source agent frameworks and libraries along with evaluation may start to gain more traction.


The technology, media, and telecom and advanced industries sectors are leading the way in use of open source AI. Among individual industries with large numbers of respondents (more than 70 percent), the technology industry is in the vanguard, with 72 percent of respondents using open source models and 39 percent of them having used open source for more than three years. In terms of geography, India, the United Kingdom, and the United States lead on open source model use, likely reflecting the relative maturity of their technology industries (Exhibit 3). Regardless of industry or location, respondents characterizing AI as important or very important to their organization's competitive advantage (81 percent of organizations) were more than 40 percent more likely to report using open source AI models, tools, and data—suggesting that organizations with more AI maturity, and more experienced developers, are more likely to be using open source AI technologies.

Mark Surman, president, Mozilla Foundation: “The momentum behind open source AI is undeniable. In just the past year, we’ve seen countless examples proving that community-driven innovation can not only compete with but even outperform proprietary models. The next big bet is building open tools and a stack that make AI truly accessible—like an AI Lego box that anyone can use. If we get this right, open source AI won’t just be an alternative to closed systems. It will be the foundation for a more competitive, creative, and innovative future.”

Exhibit 3

India and the tech sectors are hot spots for regular use of adopting open source AI models.

Organizations’ regular use of open source AI models, by region and industry,¹ % of respondents

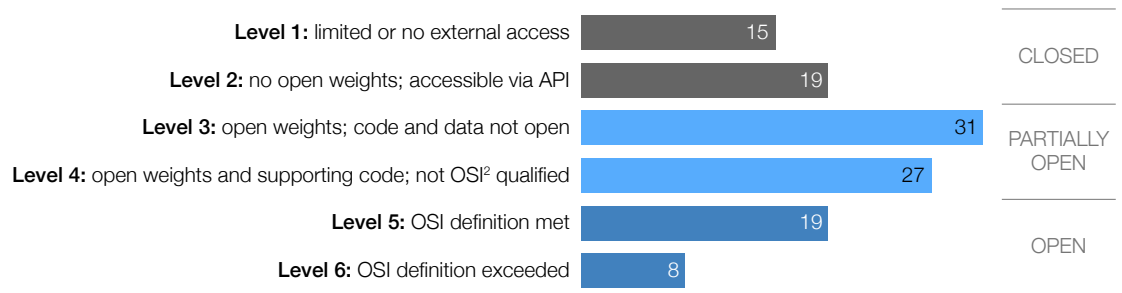


¹Advanced industries (including semiconductors), n = 52; technology, media, and telecommunications, n = 260; financial and professional services, n = 157; energy and materials, n = 43; public sector, n = 59; healthcare and life sciences, n = 53; consumer, n = 45; other, n = 34.
Source: McKinsey Open Source AI Survey, 703 participants with experience in working with AI tech systems, Dec 9, 2024–Jan 24, 2025

For AI models, respondents were asked about their preferences across the openness scale. Respondents most frequently selected partially open models—that is, models with open weights but that do not necessarily meet the OSI standard for openness with public data set information (Exhibit 4). This finding likely reflects the current competitive landscape, in which partially open models include many of the best-known and well-resourced models, such as open weights models like the Llama 3 and 4 families, and models with commercial-usage restrictions (such as Stable Diffusion).

Most organizations prefer to use AI models classified as partially open.

Interest in using AI models in production, by level of access openness, % of respondents¹



¹Respondents could select >1 response.

²Open Source Initiative.

Source: McKinsey Open Source AI Survey, 703 participants with experience in working with AI tech systems, Dec 9, 2024–Jan 24, 2025

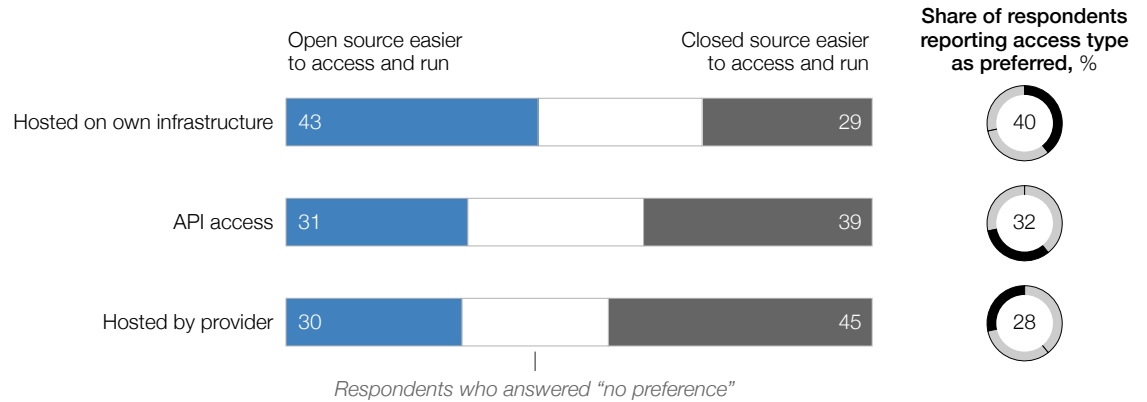
Another reason for the appeal of open weights models may be that such models can be hosted on an organization’s own infrastructure, a preference cited by 40 percent of enterprise leaders. Leaders interviewed indicated that self-hosted LLMs offer more control over data privacy, reducing the risk of data breaches involving third-party vendors and allowing for the implementation of customized security protocols—an approach that minimizes external access to sensitive data and gives full control over data management and security protocols.

For models hosted on organizational infrastructure, 43 percent of respondents believe open source offerings are easier to access and run (Exhibit 5). Yet when models are hosted by a hosting provider or used via an API (such as Claude API, Gemini API, or OpenAI API), more respondents report that closed models are easier to access and run (45 percent and 39 percent, respectively), likely reflecting the availability and widespread use of popular APIs from leading AI companies. A distinction can be drawn between AI-forward organizations that have the skills and resources to manage their own model deployments and those willing to use an external API for the benefits of a developer suite and faster time to value. Which of these two approaches an organization adopts is often determined by several factors, such as depth of technology talent, priority use cases, and security concerns.

Lareina Yee, McKinsey Global Institute director and senior partner: “Many organizations are interested in a multimodel approach (mixing more open models with proprietary solutions). We see this evolving similarly to how many large organizations use solutions from multiple cloud service providers.”

With in-house infrastructure, organizations find it easier to run open source AI models than to run those that are closed.

Perceived ease of use and preference of AI models, by access type, % of respondents



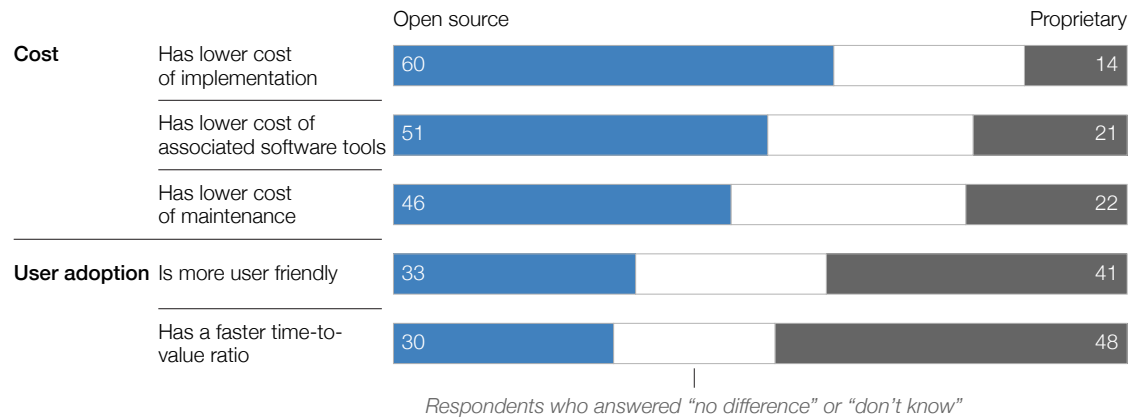
Source: McKinsey Open Source AI Survey, 703 participants with experience in working with AI tech systems, Dec 9, 2024–Jan 24, 2025

Organizations are realizing significant value from open source technologies

Many organizations using open source AI are already realizing value, and most respondents using both open source and proprietary AI report being satisfied with their AI tools. The key drivers of satisfaction are performance and ease of use. And we found that far more people are satisfied than dissatisfied with their use of open source technology—more than ten times as many respondents reported being “somewhat satisfied,” “satisfied,” or “very satisfied” with open source than the number of respondents who selected any of the “dissatisfied” options. Those using open source AI technologies report lower implementation costs (60 percent of decision-makers) and lower maintenance costs (46 percent of decision-makers) compared with proprietary tools. Fifty-one percent of respondents believe using open source lowers the costs of using associated software tools in respondents’ organizations (Exhibit 6).

Tech leaders say that open source AI models are less costly to deploy than proprietary AI models are but lag behind them in time to value.

Cost and user adoption of open source vs proprietary AI models, % of respondents



Source: McKinsey Open Source AI Survey, 703 participants with experience in working with AI tech systems, Dec 9, 2024–Jan 24, 2025

Open source AI use is more common in large organizations and industries already investing in AI, since they are more likely to have the resources to invest in talent and internal support. Similarly, providers of proprietary AI models have made significant investments in ease-of-use and developer tools, aiming to deliver a seamless experience. This is particularly evident in consumer-focused products, where the primary incentive is to broaden adoption by offering well-documented APIs, ready-to-use code snippets, and an intuitive developer interface.

Users of both open source and proprietary tools report similar levels of value realization (in terms of revenue increases or cost savings) across business functions. The business functions that most frequently report using open source AI include IT, software engineering, and product/service development, which is the same as for proprietary AI. However, open source AI showed a slight edge in cost savings, averaging 4 percent higher than proprietary solutions, with a typical cost improvement of 26 percent. Marketing experienced the greatest cost savings with open source AI, outperforming proprietary AI by 5 percent. Service operations saw the highest revenue increases from open source AI, with a 6 percent higher revenue growth compared with proprietary solutions.


Developers report being particularly interested in open source tools—though use often reflects the experience level of the developer, since proprietary tools tend to require less expertise. Regular usage of open source AI models was more common among developers who had more experience working with AI systems (more than 40 percent more likely) than those who did not (Exhibit 7).

When asked about familiarity and comfort with open source tools, 81 percent of developers report that experience with open source tools is highly valued in their field, with 42 percent saying that open source tools are more valued than proprietary tools. In addition, a majority of developers (66 percent) report that experience with open source tools is important or very important to their job satisfaction (Exhibit 8).

Exhibit 7

The more experienced the developer, the more likely they are to regularly use open source AI models.

Regular use of open source AI models, by level of experience in working with AI systems, % of developer respondents (n = 312)



Source: McKinsey Open Source AI Survey, 703 participants with experience in working with AI tech systems, Dec 9, 2024–Jan 24, 2025

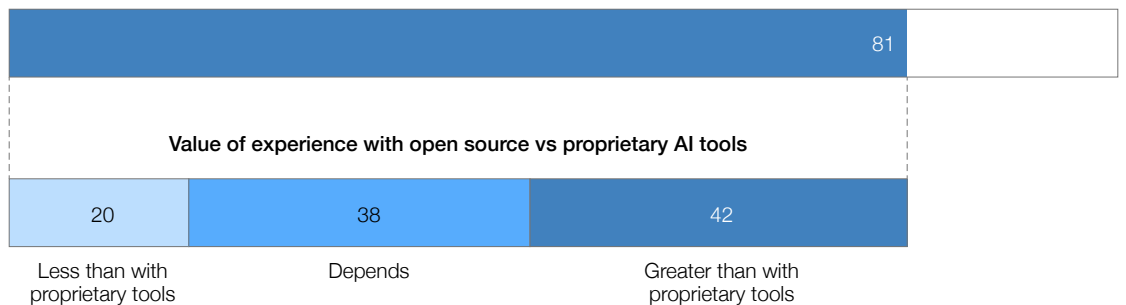
Exhibit 8

Developers say that working with open source AI tools is highly valued and important to job satisfaction.

Importance of open source AI tools to job satisfaction, % of developer respondents (n = 312)



Experience with open source AI tools perceived as valued in field, % of developer respondents (n = 312)



Source: McKinsey Open Source AI Survey, 703 participants with experience in working with AI tech systems, Dec 9, 2024–Jan 24, 2025

Harrison Chase, cofounder and CEO of LangChain: “Open source tools build up an ecosystem of integrations, which allows faster innovation through shared contributions, an ease of getting started with a wide assortment of tools, and interoperability with vector stores and model providers. As more developers adopt open source frameworks, it naturally increases the use of commercial tools as well. We see this trend continuing, with open source tools driving the development and adoption of paid services that simplify open source deployment.”

The road ahead: More experimentation, more use

The open source AI model landscape saw significant growth in 2024, marked by increased releases and improved performance parity with proprietary counterparts. If the trend continues, more models likely will be made open source. Additionally, open source initiatives such as EleutherAI, Hugging Face, and OpenMined gained significant traction, further fueling the momentum of accessible, community-driven AI development. Major milestones in 2024 included Meta’s Llama 3 (8B/70B parameters), which outperformed closed models including Claude 3 Sonnet and Gemini Pro 1.5 in benchmarks, and DeepSeek-V3, an open source model rivaling top proprietary systems in inference speed.⁴ Companies such as Apple (OpenELM) and Microsoft (Phi-3-mini) expanded open source offerings, while start-ups such as Reka AI introduced multimodal models matching GPT-4 capabilities. The past couple of years also saw collaborative efforts, such as Hugging Face’s SmolLM2 and SmolVLM, emphasizing accessibility and efficiency. While open source models lead with transparency and community-driven innovation, challenges remain in scaling training infrastructure to match proprietary systems’ training compute coupled with rapid iteration cycles.⁵

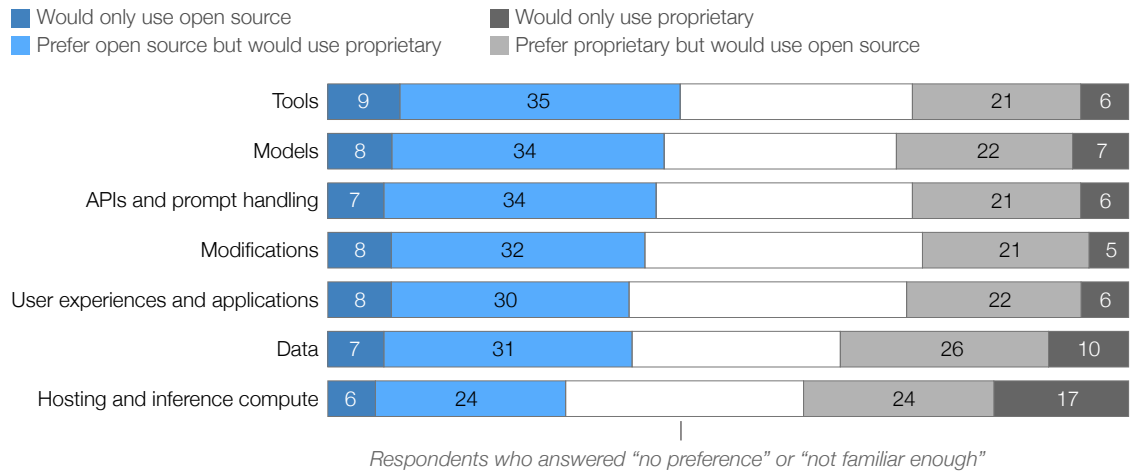
Still, on average, 75 percent of respondents expect their organizations to increase use of open source AI technologies over the next several years. Nuances will, of course, vary across each organization, and there is a range of leadership perceptions on the risks and the kinds of use cases best suited for open source. While about 70 percent of users prefer either open source or proprietary technologies across the tech stack, most indicated that they would consider both (Exhibit 9). Very few were “purists” for either proprietary or open source.

⁴ Artificial intelligence timeline, 2022 – present, AI Timeline project, GitHub, accessed March 2025; *Simon Willison’s Weblog*, “Timeline of AI model releases in 2024,” December 31, 2024; Open LLMs for code, GitHub, accessed March 2025.

⁵ Ben Cottier et al., *How far behind are open models?*, Epoch AI, November 4, 2024.

Organizations are more likely to prefer open source over proprietary AI tools across their tech stacks.

Preference for open source vs proprietary AI tools, % of respondents familiar with open source AI¹




¹Tools, n = 633; models, n = 679; APIs and prompt handling, n = 644; modifications, n = 558; user experiences and applications, n = 670; data, n = 624; hosting and inference compute, n = 645.
Source: McKinsey Open Source AI Survey, 703 participants with experience in working with AI tech systems, Dec 9, 2024–Jan 24, 2025

Looking forward, leaders' preferences are divided regarding whether they would use open source or proprietary AI in the future, with 32 percent of respondents saying their leaders have stated a preference for proprietary AI technologies, 32 percent reporting no preference, and 29 percent reporting a preference for open source. Those opting for open source AI cite cost as a key reason 63 percent of the time. Leaders with a strategic preference for proprietary AI cite security, risk, and control over the system as a top reason 72 percent of the time (Exhibit 10). When comparing similar open source and proprietary tools, respondents primarily consider security, risk, and control over the system (56 percent of respondents); cost (47 percent) and quality (46 percent) are also considerations.

Despite these reported benefits, the perceived cost differences between open source and proprietary AI may be more pronounced than the actual impact. Looking ahead, as the ecosystem of services supporting open source software becomes even more global and affordable, we anticipate that the cost savings gap between open and closed source technologies may widen further.


Leaders are drawn to the cost of open source AI technologies and the security protocols of proprietary AI technologies.

Stated preference of organization leadership for use of open source vs proprietary AI technologies, % of respondents




Reasons for preference, % of respondents selecting as top 3 response

Open source



Proprietary



Source: McKinsey Open Source AI Survey, 703 participants with experience in working with AI tech systems, Dec 9, 2024–Jan 24, 2025

While respondents are enthusiastic about the prospect of increasing their use of open source technologies, few are contributing work to the open source ecosystem. Only 13 percent of respondents indicate that they have contributed to open source projects, and 50 percent report that they are unsure about participating in future open source AI projects. For those who have participated, the most common reasons cited for doing so were talent attraction and positive brand perception. Low contribution rates may stem from the fact that open source contributions are often driven by a sense of altruism and belief in the ideals of the open source movement, rather than by direct financial benefit. Contributions from multiple stakeholders, including enterprises, academic institutions, and other foundations, are essential to drive technical progress across each area of the AI technology stack and support safe deployment and risk mitigation (see sidebar “The impact of open source on the future of model development”).

The impact of open source on the future of model development

Open source models can accelerate innovation through collaboration, reducing redundant development and fostering collective progress. Open source AI innovations are likely to have downstream impacts on two key AI technology trends: privacy-centric edge applications powered by small language models (SLMs) and the emergence of reasoning models with higher inference-time compute.

In the case of applications powered by SLMs, expect to see the emergence of industry-specific, cost-efficient language models tailored for specialized tasks and distilled into domain-specific tools to power applications. Hyperscalers such as Amazon Web Services (AWS), Google Cloud, and Microsoft Azure are already releasing such models for sectors including manufacturing and finance. Open source developers are also playing an important role in creating these SLMs, enabling the distillation process of general-purpose large language models (LLMs) with models that can match or even exceed the performance of larger ones.¹ Small models also enable edge applications and on-device intelligence for organizations that prioritize latency and/or privacy.² Some

examples of small-model hubs that distribute open source (and other) models include the Qualcomm AI Hub, which addresses the needs of edge AI product OEMs, and Ollama, which offers a framework and tools to deploy open models to the PCs of individual advanced users. We expect hubs to add trusted third-party evaluation/certification tools (for example, AILuminate from MLCommons), enhancing customer trust and confidence when selecting models.

The second key trend is the emergence of reasoning models, which employ higher compute during inference time (rather than in their pretraining time) to excel at specific tasks. While the initial wave of reasoning models were proprietary (such as OpenAI's o1 reasoning model), open source alternatives—including DeepSeek-R1 and a similarly capable model from Alibaba—have quickly followed. Other players are building on and adapting these. Perplexity has modified a version of DeepSeek³ to provide more unbiased and accurate information. Smolagents from Hugging Face has also created an alternative Deep Research model,⁴ challenging offerings

from OpenAI and Google DeepMind. Other open technologies are emerging to help builders optimize and enhance their model-training pipelines and processes. DeepSeek, for example, has continued to offer open source repositories, including parallelism and integration capabilities, for its reasoning models.

While the capabilities of open source models once lagged behind proprietary ones, base models have improved significantly. And while enterprises may face challenges in tailoring some of the components of reasoning models, the bottom line is that open source offerings now allow model service providers to bring together a full stack of technologies that delivers an effective developer experience, enables modularity, and captures the advantages of community-based development. Additionally, it provides organizations greater flexibility and choice to deploy AI either on the edge or in the cloud, depending on their privacy, latency, and performance needs. This operating model and architectural flexibility can help build more resilient AI systems, a quality especially valuable in a fast-shifting world.

¹ "How open-source is shaping the future of innovation," DevOps Online, accessed March 2025.

² Shreyas Subramanian, Vikram Elango, and Mecit Gungor, *Small language models (SLMs) can still pack a punch: A survey*, Cornell University working paper, January 3, 2025.

³ *Perplexity blog*, "Open-sourcing RI 1776," February 18, 2025.


⁴ Hugging Face Smolagents, Open Deep Research, GitHub, accessed March 2025.

Navigating the risks presented by open source AI

Amid the benefits and value of open source AI, there are a number of risks—primarily related to security and privacy—associated with open source tools that could affect their adoption. Survey respondents consider open source AI tools to be riskier than proprietary AI for most types of AI risks. The most relevant AI risks cited include cybersecurity (62 percent of respondents), regulatory compliance (54 percent), and intellectual property (50 percent). However, the risk perception varies across countries, possibly based on their regulatory environment, risk tolerance, and AI maturity (Exhibit 11). Respondents from Brazil are 10 percent more likely to identify cybersecurity as a significant risk compared with the global average. In India, 70 percent of respondents express concern about regulatory compliance, which is 31 percent higher than the average. Additionally, India shows 26 percent greater concern than other countries regarding intellectual property infringement.

Perception of AI risk varies according to country.

Perceived relevance of AI risk, by region, % of respondents



Source: McKinsey Open Source AI Survey, 703 participants with experience in working with AI tech systems, Dec 9, 2024–Jan 24, 2025

While most developers we surveyed consider open source AI to be risky, risk perception also varies according to developer experience (see sidebar “Taking action against potential risks”). More experienced developers (defined as those who have contributed to six or more AI systems in production) are far more comfortable with open AI. Such developers are about 11 percent less likely to say open source is riskier for intellectual property infringement, 15 percent less likely to say open source AI is riskier for cybersecurity, and about 11 percent less likely to say open source AI is riskier for regulatory compliance.

Taking action against potential risks

Our survey shows that enterprises perceive greater risks from open source. So what actions should they take? We see four key areas that leaders must consider when implementing a model-based system, whether open source or proprietary:

- *Guardrails:* The establishment of robust guardrails—such as automated content filtering, input/output validation, and human oversight—can help ensure responsible use and secure outputs. Open source examples include Nvidia’s NeMo Guardrails, Llama Guard, and Guardrails AI, which can aid in compliance with regulatory and ethical standards.¹
- *Third-party evaluations:* Another way to build confidence that open models will

not cause unintended harm is to conduct regular assessments with standardized benchmarks that allow for certification. During such benchmarking, private evaluations assure that test data sets are kept private from the model.²

- *Documentation and monitoring:* Operationally, a software bill of materials can help track version discrepancies and vulnerabilities by maintaining detailed inventories of open source components. Quantitative risk assessments, such as the Common Vulnerability Scoring System Calculator (CVSS) v3.0, can assess the severity of vulnerabilities in open source systems.³
- *Cybersecurity practices:* To secure data privacy and system integrity,

running models in trusted execution environments (TEEs) may help to ensure sensitive data remains encrypted during processing. Incorporating differential privacy and federated learning techniques during training can prevent models from memorizing confidential information. Strong access controls within model repositories, network segmentation between training and inference servers, continuous monitoring of security incidents, and cryptographic hash verification to confirm that models are from trusted repositories are some examples that can further address both content safety and cybersecurity challenges in production AI environments.

¹ Yi Dong et al., “Safeguarding large language models: A survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, May 2024.

² Tanmay Rajore et al., *TRUCE: Private benchmarking to prevent contamination and improve comparative evaluation of LLMs*, Cornell University working paper, June 2024.

³ Zoë Brammer et al., *Castles built on sand: Towards securing the open-source software system*, Institute for Security and Technology, April 2023; Lucie-Aimée Kaffee, “Reports on the hub: A first look at self-governance in open source AI development,” Hugging Face, June 12, 2024.


Philip Reiner, CEO of the Institute for Security and Technology: “Organizations need to return to first principles and focus on basic cybersecurity practices when it comes to AI. The reality is that open source tools often suffer from poor maintenance and outdated software. To manage this risk, companies need to regularly check compliance failure reports and assign a dedicated team to stay on top of risks and updates for any open source tools they use. Many in the C-suite don’t fully understand the risks involved, and we’re likely to see more roles shift toward quality assurance as a result.”

For respondents who do not currently use open source AI, the primary barrier to adoption is security and compliance concerns (56 percent of respondents) among others listed in Exhibit 12. However, respondents' organizations that use open source AI have begun to put technical safeguards in place during usage to mitigate some of these potential risks. More than a fifth (21 percent) of respondents using open source AI are implementing aligned weights, 35 percent are using programmable guardrails (such as Nvidia's NeMo Guardrails), 47 percent are using safeguard models (for instance, Llama Guard), and 49 percent are using prompt adjustments. Of those that are concerned or very concerned about the risks associated with training AI models, 49 percent are consulting legal counsel, 43 percent are seeking cleaned versions of models, 57 percent have technical safeguards/tests, and 36 percent are avoiding open source models altogether. Of respondents addressing copyright concerns when working with large-scale data sets, 53 percent are seeking legal counsel and 46 percent are implementing technical safeguards and testing measures. Additionally, 29 percent are purchasing cleaned data sets and 33 percent avoid open source data sets altogether.

Exhibit 12

Leading barriers to adopting open source AI tools are concerns about security and compliance, long-term support, and intellectual property.

Reported barriers to adoption of open source AI tools, % of respondents



Source: McKinsey Open Source AI Survey, 703 participants with experience in working with AI tech systems, Dec 9, 2024–Jan 24, 2025

The future of open source AI

Open source AI is becoming a key part of the emerging AI landscape. Our survey shows a strong and growing demand for open source technologies across industries, geographies, and the tech stack. As the open source ecosystem expands and use increases, organizations will have a significant opportunity to reduce development costs and deploy customized AI systems on edge devices. However, to fully harness the benefits of open source technology, businesses may need to have higher confidence that they can address the technical and legal risks that they perceive to be greater for some open source AI usage.

Additionally, since experienced developers are the primary contributors to and users of open source projects, upskilling teams is essential. Increased collaboration in this space is also likely to drive the emergence of new services aimed at enhancing the developer experience. Much like we've observed in the cloud and software industries, a hybrid approach will likely become the standard, with open source and proprietary technologies coexisting across multiple layers of the AI technology stack to meet diverse organizational needs.

About the authors and acknowledgments

Ankit Bisht is a partner in McKinsey's Dubai office, and **Lareina Yee** is a senior partner in the Bay Area office, where **Roger Roberts** is a partner, **Brittany Presten** is an associate partner, and **Katherine Ottenbreit** is a consultant.

The authors wish to thank their research partners at the Mozilla Foundation and the Patrick J. McGovern Foundation; their colleagues in QuantumBlack Labs, the software development and R&D arm of QuantumBlack, AI by McKinsey, which brings AI innovations to clients and has made many contributions to the open source software ecosystem in AI and machine learning; Cayla Volandes in McKinsey's New York office; Natasha Maniar in McKinsey's Bay Area office; and their external academic collaborators for their insights and perspectives on the survey draft and analysis, including Knut Blind at Fraunhofer ISI, Luca Vendraminelli at Stanford University, Sayash Kapoor at Princeton University, and Genevieve Smith at University of California, Berkeley.

Appendix

Respondent demographics

Our survey ran from December 2024 to January 2025, covering 703 respondents in 41 countries, with a particular focus on Brazil, France, India, the United Kingdom, and the United States. See a summary of respondent demographics below.

— *Geographies:*

- United States: 150
- United Kingdom: 143
- India: 115
- France: 84
- Brazil: 73
- Europe (excluding France): 70
- Other: 68

— *Industries:*

- Advanced industries (includes advanced electronics; aerospace and defense; automotive and assembly; semiconductors): 52
- Consumer (includes consumer and packaged goods; retail): 45
- Energy and materials (includes electric power and natural gas; engineering, construction, and building materials; oil and gas; metals and mining): 43

- Financial and professional services (includes business, legal, and professional services; financial services; private equity and principal investors): 157
- Healthcare and life sciences (includes healthcare; pharmaceuticals and medical products): 53
- Public sector (includes public sector; social sector): 59
- Technology, media, and telecommunications (includes media and entertainment; technology; telecommunications): 260
- Other (includes agriculture; chemicals; real estate; travel, logistics, and infrastructure): 34

Throughout the survey, we categorize respondents into two main archetypes: decision-makers and developers. Decision-makers are executives, vice president–level and above. Developers are senior developers/technologists or other equivalent roles. Respondents were required to have experience in at least one area of the AI technology stack. Developers were also screened out if they had not contributed to AI systems

in production. See the breakdown of respondent job titles below.

— *Decision-makers: 391 respondents*

- Titles with at least five respondents: chief analytics/science officer, chief data officer, chief digital officer, chief executive officer, chief financial officer, chief information officer, chief operating officer, chief strategy officer, chief technology officer, department/division/business unit head, general manager, managing director, other C-suite officer, other executive (vice president–level or above), owner/partner/principal, senior vice president, vice president

— *Technologists: 312*

- Titles with at least five respondents: lead engineer (or technologist equivalent), other senior technologist/engineer (senior developer/architect level and above), principal engineer (or technologist equivalent), senior engineer (or technologist equivalent), senior staff engineer (or technologist equivalent), staff engineer (or technologist equivalent)

April 2025
Copyright © McKinsey & Company
Designed by McKinsey Global Publishing

[McKinsey.com](https://www.mckinsey.com)

