

mlinsider 2023 Survey

Full Report

The state of Generative AI and
Machine Learning at the end of 2023

by cnvrg.io
an Intel Company

Table of Contents

What is ML Insider? _____	4
Who took this survey? _____	4
About cnvrg.io, an Intel company _____	4
Key takeaways and trends _____	5
Section 1: Survey demographics _____	8
Section 2: AI/ML adoption and maturity _____	11
Section 3: Generative AI and Large Language Models _____	16
Section 4: MLOps and infrastructures _____	22
Section 5: AI transparency, ethics and challenges _____	27

What is ML Insider?

The ML Insider is an annual analysis by cnvrg.io, an Intel company, of the Machine Learning industry highlighting key trends, points of interest, and challenges that AI developers experience every day. This report offers insights into how over 434 AI professionals are building, training, deploying and adapting their machine learning stack to better suit today's modern, complex ML and Large Language Model (LLM) workflow.

Who Took This Survey?

The survey covered 434 participants, representing organizations in size from a few employees to over 5,000. 51% of respondents came from companies with 600 employees or larger. The insights from the survey covered dozens of industries; among the most common were: Information Technology Services and Computer Software and Financial Services/Banking. Education, Healthcare, Consumer Goods, Telecommunications, Automotive, Insurance Media/Entertainment, Defense and those that selected "Other" represent the other half of respondents. 28% of all respondents are data scientists, 23% engineering/DevOps and 17% in software development roles. More detailed demographic information can be found in section 5 of the report.

About cnvrg.io, an Intel company

cnvrg.io, an Intel company, is a full stack machine learning operating system with everything an AI developer needs to build and deploy AI on any infrastructure. cnvrg.io was built by data scientists, to help data scientists and developers automate training and deployment of machine learning pipelines at scale and help organizations enhance their products with GenAI and LLMs by making it more cost effective and easier to deploy large language models on Intel's purpose built hardware provided by Intel Developer Cloud. cnvrg.io helps organizations accelerate value from data science, and leverage the power of generative AI technology faster.

Key Takeaways and Trends

- **AI integration remains low as this continues to be a difficult area to execute**
58% of respondents currently have a low level of AI integration with no more than 5 models running. Only 11% are running over 50 ML models. These results are statistically equal to prior years. 26% of large companies with more than 5,000 employees are running 50 or more ML models in their organization.
- **62% of AI professionals continue to say it is difficult to execute a successful AI project**
10% consider it extremely difficult while 52% said it is just difficult. The survey also reveals that the larger the company, the more difficult it is to execute a successful AI project.
- **Most organizations realize the benefits from AI solutions, but substantially fewer than last year**
75% indicated that their organization will realize the benefits of their AI solutions in terms of ROI, productivity and optimization. While this represents the majority, it also represents a significant decrease from last year's 89%. 83% of large companies (more than 5,000 employees) have realized benefits compared with only 65% of small companies with no more than 10 employees.
- **Lack of expertise, budget and finding AI talent are the top challenges organizations are facing when it comes to executing ML programs**
Cost is a primary challenge for smaller companies under 200 employees. While companies of all sizes find lack of knowledge and expertise to be a big challenge, larger companies find knowledge and expertise to be a top challenge, along with hiring talent.
- **Only 25% of organizations have deployed Generative AI models to production in the past year**
U.S. based respondents are significantly more likely than those outside the U.S. to have deployed Generative AI models to production in the past year (40% vs 23% of those outside the U.S.).
- **Of those who have deployed Generative AI models in the past year, several benefits have been realized.**
About half said they have seen better customer experiences (58%), improved efficiency (53%), enhanced product capabilities (52%), and cost savings (47%).

- **Nearly half of AI professionals see infrastructure as the largest barrier to productionizing LLMs**
Infrastructure (46%) is the biggest challenge in productionizing LLMs, mentioned twice as much as the 2nd biggest challenge – monitoring (21%).
- **Concerns exist with implementing LLMs**
The greatest concerns with implementing LLMs is compliance & privacy, reliability, and the high cost of implementation, and a lack of technical skills.
- **8 out of 10 respondents admit their skills need to improve due to the increased interest in LLM adoption**
These respondents are split between having what they need to improve these skills and not knowing where to start.
- **AI teams spend an average of 145 days from experimentation to production**
57% get from experimentation to production within 90 days. Respondents that were data scientists reported that they only need 109 days to get from experimentation to production.
- **Teams are less concerned with explainable AI this year and are more likely to have explainable AI as a primary KPI in their organization**
30% are not concerned with addressing explainable AI (up from 20% compared to 2022). Fewer are planning to introduce explainable AI techniques in the next 12 months and fewer are also using specific AI explainability tools.
- **The complexity and lack of AI talent are the greatest obstacles to AI adoption and acceptance**
53% consider the technical complexity of AI development to be the greatest obstacle to universal AI adoption and acceptance. 46% said the lack of AI talent is the top obstacle. AI governance and cost of development and maintaining AI are also concerns for 44% of survey respondents. Larger companies are more likely to be concerned with technical complexity compared to smaller companies. (just 35% compared with more than 50% of larger companies).




ML Insider 2023
Results and Analysis


● Demographic information

The following reports the demographics of the survey respondents.


● Primary roles



● Team size




Salaries




● > \$200k/yr ● \$120K-\$200K ● \$75K-\$120K ● \$50K-\$75K ● <\$50k/yr ● Rather not say

Industries



Company size



Gender



● AI/ML adoption and maturity

Overall, AI maturity remains quite low. In this section, we measure AI maturity by the number of models an organization has running in their product. Note that organizations may in fact have AI that is not in production or applied to a product solution.

● Figure 1: Who is involved in building AI solutions in your organization?

While data scientists have historically been the primary builders of AI technology, over the past 2 years the survey has revealed that developers and engineers are taking on a larger role in the AI process. There is also more involvement from DevOps, Analysts, and business users compared to what we have seen in previous years.




Figure 2: What level of AI integration exists in your organizations applications?

2022 vs 2023 comparison

Since the 2022 ML Insider survey, AI has not seen the acceleration of maturity that was expected. Despite the large investments predicted to be made in AI in 2023, shockingly 58% of respondents report that they have 0-4 models running in their products, 1% higher than in 2022 with 57% reporting 0-4 models in their products.

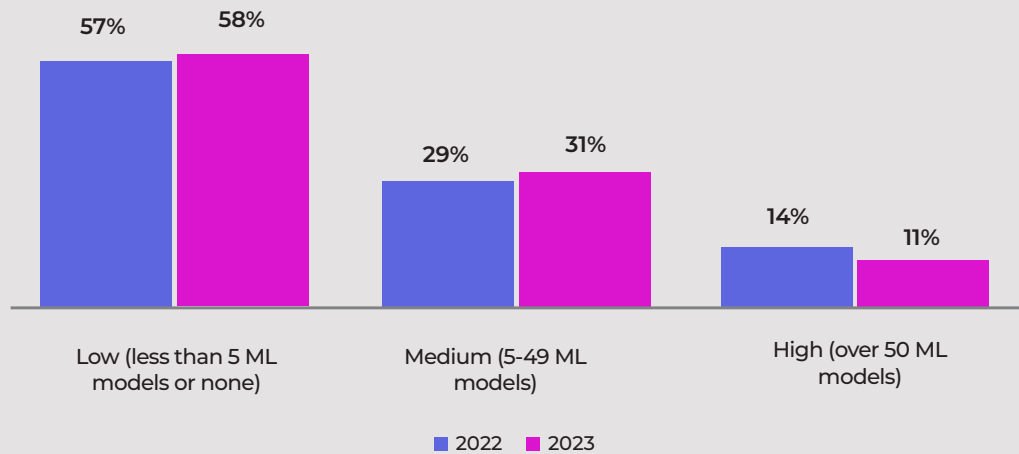


Figure 3: Level of AI maturity by company size

The larger the company size correlates with a higher level of AI adoption, whereas smaller companies are having low to no models running in production. 25% of large companies with more than 5,000 employees are using 50 or more ML models

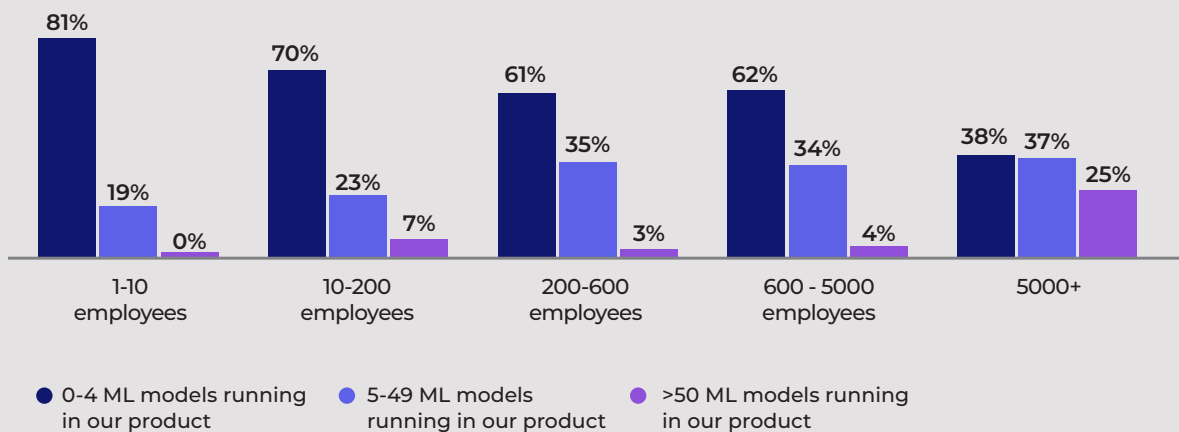


Figure 4: Level of AI maturity by industry

Industries that have the highest level of AI adoption include Financial Services/Banking, Defense, Insurance, and IT. Industries with the lowest level of AI adoption are Education/E-learning, Automotive, and Telecommunications which is consistent with 2022 responses.




Figure 5: How would you rate the difficulty of executing successful AI projects?

Two-thirds of AI professionals continue to say it is difficult to execute a successful AI project.




Figure 6: Level of difficulty by company size

The larger the company, the more difficult it is to execute a successful AI project (only 52% of small companies with no more than 10 employees but 67% of large companies with more than 5,000 employees).




Figure 7: AI/ML use cases

2022 vs 2023 comparison

Compared to 2022, the use of chatbots/virtual agents, as well as translation/text generation have seen a spike as a popular AI use case. That could be due to the rise in Large Language Model technology in 2023 and the advance in generative AI technology.




Figure 8: Has your organization realized benefits (e.g. ROI, productivity, optimization) from your AI solutions?

2022 vs 2023 comparison

Compared to 2022, in 2023 organizations are seeing fewer benefits from their AI solutions. 75% indicated that their organization has realized the benefits of their AI solutions in terms of ROI, productivity and optimization. While this represents the majority, it also represents a significant decrease from last year’s 89%. 86% of respondents based in the U.S. have realized benefits which is significantly higher than respondents from outside the U.S. (73%).

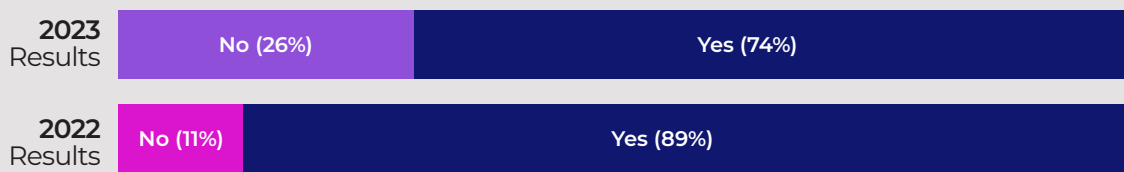



Figure 9: Which phase of the AI pipeline do you find most challenging?

2021 vs 2022 vs 2023 comparison

Lack of data is a rising problem in AI and becoming even more significant with the rising popularity of large language models. As more companies require larger amounts of data to train on, data acquisition is becoming one of the most challenging phase in the AI pipeline. Data cleansing ranked 2nd but is less of a challenge compared with prior years (20% in 2021 and 19% in 2022). Model training has become more of a challenge this year with 8% in 2022 vs 12% in 2023 reporting it as a challenge. This could be due to the increased challenge of training larger language models.




● Generative AI and Large Language Models (LLMs)

As the survey reveals, there are already subtle differences in the market since the availability of Large Language Models (LLMs). In this section we delve deeper into the impact that Generative AI (GenAI) technology and LLMs are having on the market, AI developers, and AI technology as a whole.

● **Figure 10:** How important is Generative AI in your organization's future AI strategy?

It's difficult to argue that generative AI will not play a large role in the future of AI technology. Though, surprisingly, the majority of respondents found that GenAI is only moderately important to the future of organizations strategies. 44% consider Generative AI to be highly important.



● **Figure 11:** What is your current organization's maturity level in using generative AI?

Despite the hype in generative AI, only 10% of respondents said that they had already launched GenAI solutions to production. The majority of organizations are still in the research or testing phase of incorporating generative AI in production. Additionally, when asked, only a quarter of respondents had declared that their organization had deployed any generative AI models to production in the last year.




Figure 12: Has your organization deployed any GenAI models to production in the past year?

75% of respondents reported that their organization has yet to deploy generative AI models to production in the past year. The survey also revealed that U.S.-based respondents are significantly more likely than those outside the U.S. to have deployed Generative AI models to production in the past year (40% vs 22% of those outside the U.S.)



Figure 13: GenAI deployment by industry

How are organizations approaching Large Language Models? Are they investing in building their own Large Language Models leveraging open source models, or leveraging pre-built solutions? The majority of organizations are building their own Large Language Model solutions and customizing to their use case. Many are also taking a 2 tierd approach where they are combining the build + buy strategy. As this is the early stage of LLM maturity, we may say this statistic change as more prebuilt solutions gain in availability, and LLM services rise.




Figure 14: What is your organization's approach to implementing LLM use cases?

Are more organizations developing their own Large Language Model or using pre-built solutions? The majority of organizations are building their own Large Language Model solutions and customizing to their use case. Many are also implementing a 2 tierd approach where they are combining the build + buy strategy. As this is the early stage of LLM maturity, we may say this statistic change as more prebuilt solutions gain in availability, and LLM services rise.

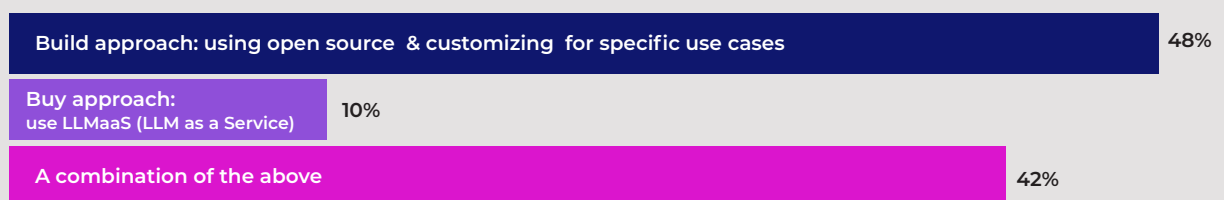


Figure 15: What is your biggest challenge productionizing LLMs?

Nearly half of AI professionals see infrastructure as the largest barrier to productionizing LLMs. Infrastructure (46%) is the biggest challenge in productionizing LLMs, mentioned twice as much as the 2nd challenge – monitoring (20%).




Figure 16: What are your top 3 LLM use cases?

LLMs are being used for a variety of use cases. Chatbots are the top LLM use case but coding/debugging, content development, customer support, text summarization, and data enrichment are use cases for about a third of respondents.




Figure 17: What benefits has your organization seen from the use of LLMs?

Organizations are seeing multiple benefits from the application of LLMs. Of those that have implemented LLMs into their organization in the past year, 27% reported using LLMs to improve customer experience, efficiency, product capabilities, and slightly fewer answered cost savings.




Figure 18: What is your greatest concern with implementing LLMs into your business?

The survey reveals a diverse set of concerns regarding the implementation of Large Language Models (LLMs) in business operations. Compliance and privacy emerged as the top worry for 28% of respondents, emphasizing the importance of adhering to regulations. Reliability closely followed at 23%, underlining the need for consistent and accurate results. High implementation costs troubled 19%, highlighting financial considerations. Additionally, 17% cited a lack of technical skills, and 10% expressed concerns about latency. A smaller 3% mentioned miscellaneous concerns. This survey offers valuable insights into the multifaceted considerations surrounding LLM integration in business.




Figure 19: LLM concerns by company size

The survey reveals that the larger the company, the greater compliance & privacy is a concern. Reliability is a concern across all company sizes.




Figure 20: How well do you understand the mechanisms of how LLMs generate responses?

Since the rise of GenAI, there has been a growing concern for trust and understanding of how large language models generate responses. Of the AI professionals surveyed, only 19% reported having a strong understanding of how LLMs generate responses. Though, impressively, 44% of respondents reported a good grasp of these mechanisms, indicating a substantial familiarity with the inner workings of LLMs.




Figure 21: Do you believe you need to improve your skills in your role due to the increasing interest in LLM adoption?

8 out of 10 respondents admit their skills need to improve due to the increased interest in LLM adoption, highlighting a strong awareness of the evolving demands in their roles. In contrast, a mere 8% expressed confidence in their current skill set, indicating that they do not perceive a pressing need for improvement.





Figure 22: Do you have the necessary tools and support to improve your skills?



Figure 23: What are your organization's top priorities for AI in the next year?

Top priorities in AI for most respondents are on LLM, and general AI innovation as well as scalability as opposed to focusing on ethics and transparency




● MLOps and Infrastructure

This section dives into the tools and infrastructures used to build and deploy AI. From MLOps platforms to compute resources, we examine the cost, benefits, and popularity of leading AI tools.

● **Figure 24:** How long does it typically take to get from experimentation (model staging, training, testing) to production in your organization?

2022 vs 2023 comparison

Survey reveals AI teams spend an average of 145 days from experimentation to production. There has been a slight improvement in AI execution in the past year.



● **Figure 25:** What does your company use for MLOps?

62% of organizations are using some form of MLOps solution whether in house developed or an MLOps platform. Though, 20% of respondents reported not using any MLOps practice in their organization, while some simply responded that they did not know.




Figure 26: Where do you primarily run your ML workloads?




Figure 27: How important is the ability to shift AI workloads between cloud vendors in your operation?



Figure 28: Which cloud providers do you use? (Select all that apply)

2022 vs 2023 comparison

It is no surprise that AWS, Azure and GCP remain the most commonly used cloud providers. Though, compared to last year, a larger percentage of respondents selected more than one cloud provider suggesting that more organizations have adopted a multi-cloud infrastructure.




Figure 29: How significant is adopting a multi-cloud environment to your overall innovation strategy?

More than half of survey respondents see multi-cloud environments as an important part of their AI strategy. Just over half (56%) said that adopting a multicloud environment to their strategy is moderately, very, or extremely important to their organization.




Figure 30: Why is your organization considering a multi-cloud infrastructure?

What is making organizations adopt or consider adopting a multi-cloud infrastructure? Cost effectiveness and reducing vendor lock-in are the top reasons for considering a multi-cloud environment.




Figure 31: What is the biggest challenge for adopting a multi-cloud environment?

The complexity and integration is the top challenge in adopting a multi-cloud environment (35%).




Figure 32: What is the average CPU count per task that you or your team use for running ML tasks?




Figure 33: What is the average GPU count per task that you or your team use for running ML tasks?

2021 vs 2022 vs 2023 comparison

CPU and memory counts have remained consistent while GPU usage is showing shifts. Only 16% are using just 1 GPU count per task which has been trending down since 2021 while those using more than one GPU per task has grown since 2021.




Figure 34: What is the average memory count per task that you or your team use from running ML tasks?

2021 vs 2022 vs 2023 comparison




Figure 35: How frequently are your datasets updated/refreshed for AI?

2022 vs 2023 comparison

Datasets are updated for AI less frequently than a year ago. 57% update or refresh their datasets at least once a month. While this represents the majority, it is a significant decrease from a year ago when 68% refreshed datasets monthly.





Figure 36: On average, how much do you spend per month on cloud costs?

Cloud costs average about \$43k per month. Internal analysis found that cloud costs among U.S. respondents is significantly higher than costs for those outside the U.S. (average of \$69k vs just \$44k monthly).



● AI Transparency, Ethics and Challenges

The 2022 survey revealed that ethics and transparency were a major concern for organizations with AI. As we understand from section 1, AI adoption has continued to see slow growth and maturity. In this section we dive in to understand the main challenges and concerns facing the future of AI.

Figure 37: What is the primary challenge your company faces in executing ML programs?

Lack of expertise, budget and finding AI talent are the top challenges organizations are facing when it comes to executing ML programs. Compared to 2022, there was a significant rise in respondents that rated lack of budget as a primary challenge.




Figure 38: Common challenges by level of maturity

Organizations with medium to high AI adoption find budget constraints and hiring talent to be amongst the top challenges. While companies with lower to no AI adoption find lack of knowledge and expertise to be their greatest challenge when implementing AI.




Figure 39: Primary challenge by company size

Cost is a primary challenge for smaller companies under 200 employees. While companies of all sizes find lack of knowledge and expertise to be a big challenge, larger companies find knowledge and expertise to be a top challenge, along with hiring talent. [Compared to 2022 data](#), where budget was not a primary challenge, we are seeing that even larger organizations saw budget restraints for AI projects as a top challenge in 2023 which could be a result of the economic situation in 2022 vs 2023.

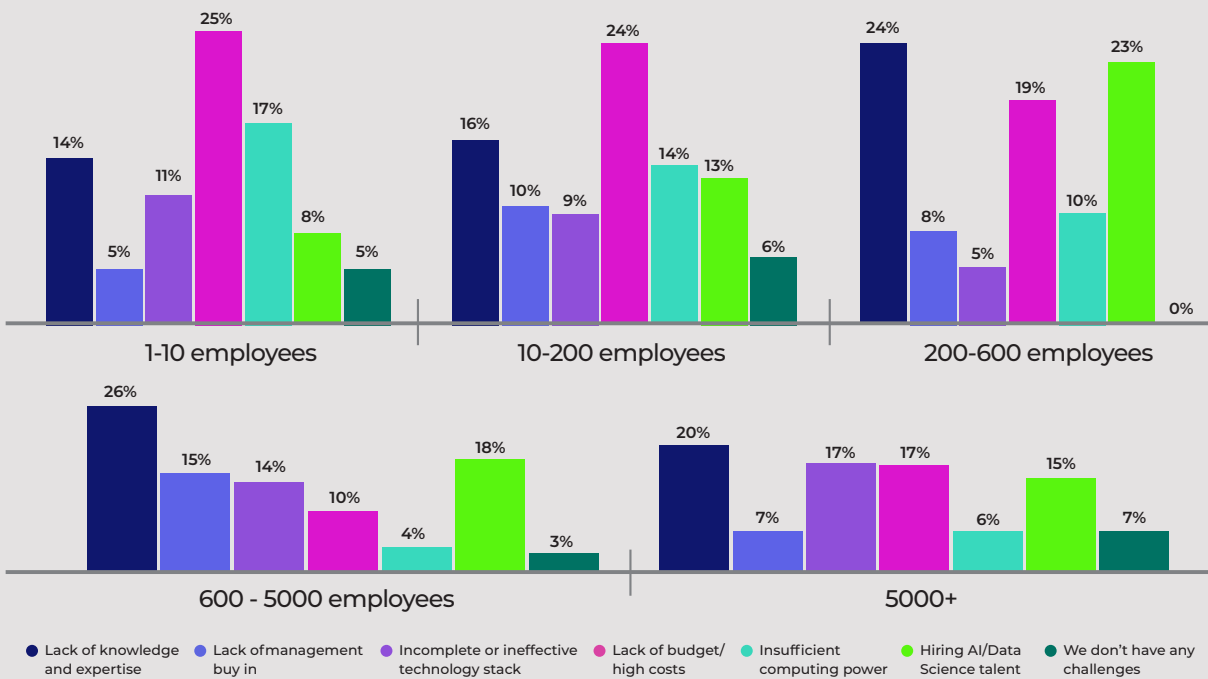


Figure 40: How is your team addressing explainable AI?

2022 vs 2023 comparison

At what level are organizations addressing AI explainability? Teams are less concerned with explainable AI this year and are more likely to have explainable AI as a primary KPI in their organization. 30% are not concerned with addressing explainable AI (up from 20% a year ago). Fewer are planning to introduce explainable AI techniques in the next 12 months and fewer are also using specific AI explainability tools.





Figure 41: In your opinion, what is the greatest obstacle to universal AI adoption and acceptance?

The majority of respondents believe the technical complexity of AI development will be the biggest challenge to universal AI adoption and acceptance. Today, it is quite difficult to build AI, and there is little transparency and understanding about how these applications are being made. In order to achieve universal AI adoption and acceptance, simplification in the development process will be required.






● Summary

Overall, the 2023 ML Insider revealed that AI maturity has not changed significantly over the last few years. The rise in GenAI technology has shifted the industry in some ways, but organizations are slow to adopt it. Despite the hype, only 25% of organizations have deployed any genAI models to production in the past year. The survey found that infrastructure is ranked as the largest barrier to productionizing LLMs and that 84% of respondents admit their skills need to improve due to the increased interest in LLM adoption. The survey also revealed that U.S. based respondents (40%) are significantly more likely than those outside the U.S. (22%) to have deployed GenAI models.

The survey reveals a myriad of challenges that might be causing a slow adoption of LLM technology in businesses such as lack of knowledge, cost, as well as compliance. 84% of respondents admit that their skills need to improve due to the increasing interest in LLM adoption, while only 19% said they have a strong understanding of the mechanisms of how LLMs generate responses. This reveals a knowledge gap as a potential barrier to GenAI adoption. Additionally, respondents ranked compliance & privacy (28%), reliability (23%), the high cost of implementation (19%), and a lack of technical skills (17%) as the greatest concerns with implementing LLMs into their business. Nearly half of respondents see infrastructure as the biggest technical challenge to productionizing LLMs.

However there is no doubt that GenAI is having an impact on the industry. Compared to 2022, the use of chatbots/virtual agents has spiked 26% and translation/text generation is up 12% in 2023 as popular AI use cases. That could be due to the rise in LLM technology in 2023 and the advance in GenAI technology. Organizations that have successfully deployed GenAI in the past year have seen quite a few benefits from the application of LLMs such as better customer experience (27%), improved efficiency (25%), enhanced product capabilities (25%) and cost savings (22%).

If you are interested in learning more about how you can accelerate your AI and LLM innovation, [request a demo with cnvrg.io](#) MLOps or join our [LLM Accelerator Program](#) to help you solve your MLOps, LLM and AI hardware challenges with a complete, secure AI hardware and software solution.



Thank **you**