vertesia

The State of Enterprise LLM Preparedness and Adoption

Introduction

We are officially living in the era of Large Language Models (LLMs).

Not only are organizations now incorporating strategies to leverage AI to drive efficiency and productivity by automating time-consuming workflows, but AI-driven data analysis, along with communication, collaboration, and personalization tools are also starting to transform business operations across all sectors.

The revolutionary power of LLMs is undeniable.

However, while LLM adoption is gaining momentum, most enterprises are still in the early experimental stages of integrating these into business solutions, laying the foundations for bona-fide LLM use in the coming years.

As they explore this groundbreaking technology, tech teams are beginning to understand that relying on just one AI tool will not afford them the true benefits of integrating LLMs with their business applications.

In other words, the time for single model generative AI experimentation is over, and organizations must prepare themselves for a world in which managing multiple AI models and inference providers is the norm.

To find out how much progress has been made on this front, we conducted research among senior tech professionals in a range of job roles at companies of varying shapes and sizes.

Based on the findings, this report reveals who is most prepared to fully integrate LLMs into their business applications and what they are doing to ensure success as they test and deploy multiple models and providers.

Hopefully, the insights we share will help you as you plan your organization's AI strategy for LLM orchestration and management, no matter what stage of the LLM journey you're at.
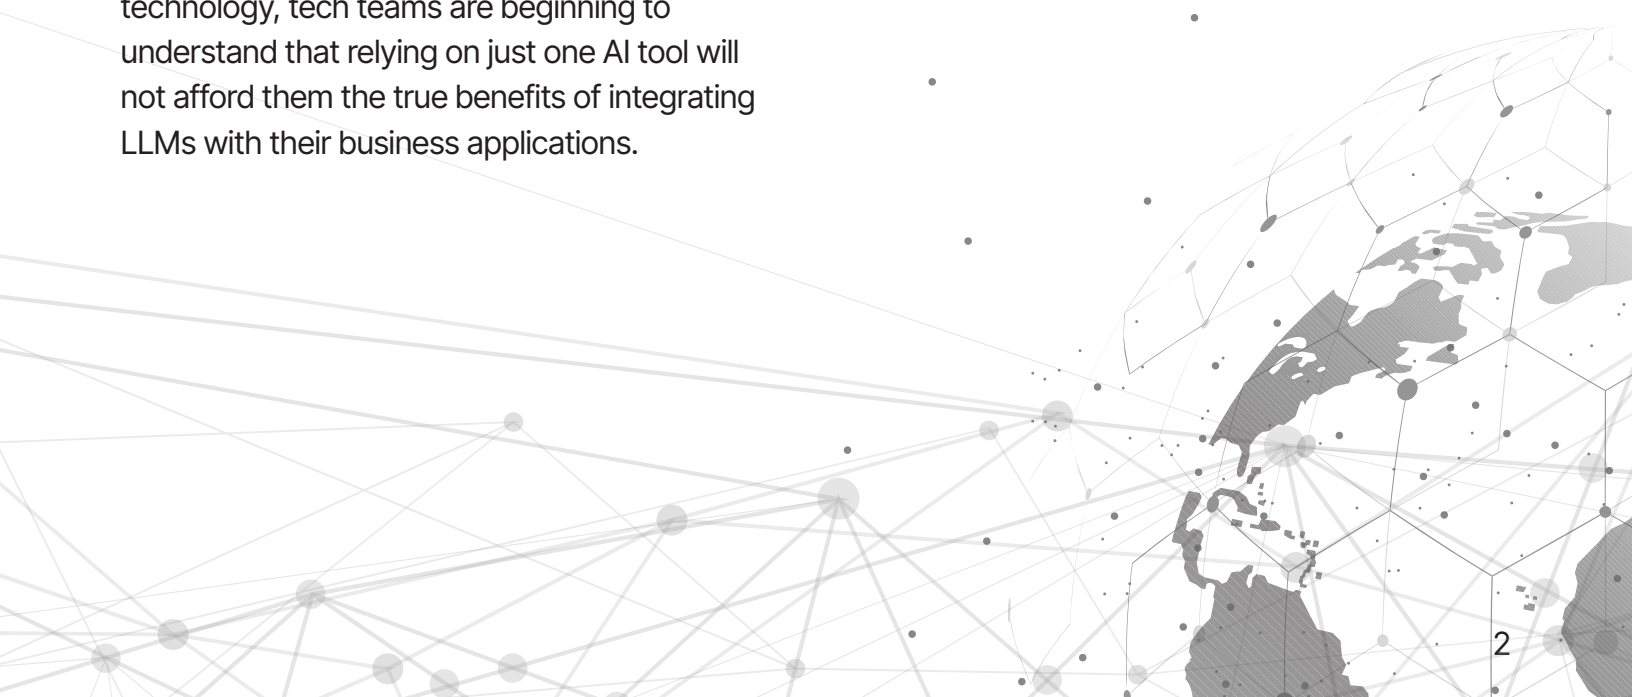
Methodology

Our research partner, Censuswide, surveyed 502 senior tech professionals at US-based companies with 1000+ employees between 26.03.2024 - 02.04.2024. Respondents held IT, Solutions, or AI/Machine Learning job roles in the banking, financial services, payer healthcare, and insurance industries.

As part of our analysis, we assessed how prepared respondents are to adopt LLMs within their organization by giving them a score based on their answers to the following initial segmenting questions:

Q1. In the next two years, how frequently do you plan to initiate Large Language Model (LLM) projects for enterprise solutions?

Q2. In the next two years, how many models would you expect to run (considering that a fine-tuned model is a different model)?

The groupings, as defined below, provide unique insights into what is impacting organizations' preparedness and their views on the future of LLM technologies.

Score	Group
Low (0-4)	Unprepared
Medium (5-6)	Somewhat prepared
High (7-9)	Very prepared

Censuswide abides by and employs members of the Market Research Society which is based on the ESOMAR principles.

The Undeniable Value of LLMs

According to the findings, senior tech professionals surveyed recognize LLMs' potential value to their organization. Of those surveyed, 9 out of 10 respondents (90%) believe fine-tuned LLMs would bring value[1] to their organization, with 42% attributing significant value to LLM usage.

90% say fine-tuned LLMs would bring value[1] to their organization

At the same time, many senior tech professionals expect that implementing a 99% accurate LLM for a specific set of tasks would yield great benefits within the first year of use.

For example, almost 3 in 5 (58%) say they would expect to materially shorten the duration of information-rich processes with automation, and over half (55%) expect LLM-powered processes to have better outcomes in terms of overall quality and speed. While operational efficiency gains from LLMs are expected to impact cost reduction (53%), almost 2 in 5 (39%) expect new products to be developed in the first 12 months of deploying a 99% accurate LLM for a specific set of tasks.

Faster Processes with Automation

Better Quality Outcomes & Speed

Cost Reduction

New Products

The question is, are organizations ready to benefit from integrating LLMs? Or is there more work to be done before businesses can achieve true transformation?

In the following sections, we'll find out how many tech teams are prepared to delve into the world of LLMs and start reaping the rewards of integrating them with their enterprise solutions, as well as revealing the challenges that are holding others back from embracing this cutting-edge technology.

Notes:
[1]'Significant value' and 'Some value' responses combined.

Who's Ready for LLMs?

To find out how prepared organizations are to adopt LLMs in a multi-provider market, we asked senior tech professionals to tell us how many models they would expect to run in the next two years.

We discovered that over 9 in 10 (94%)[1] respondents are expecting to run more than one LLM, with the highest percentage saying they expect to run 6-9 models (42%).

94% are expecting to run more than one LLM in the next 2 years[1]

These findings, along with data on how many LLM projects those we surveyed plan to initiate in the next two years, give a good indication of how far along organizations are in their preparations for integrating multiple LLMs with their enterprise systems.

By creating a scoring system based on these findings, we were able to rank respondents according to their overall levels of preparedness.

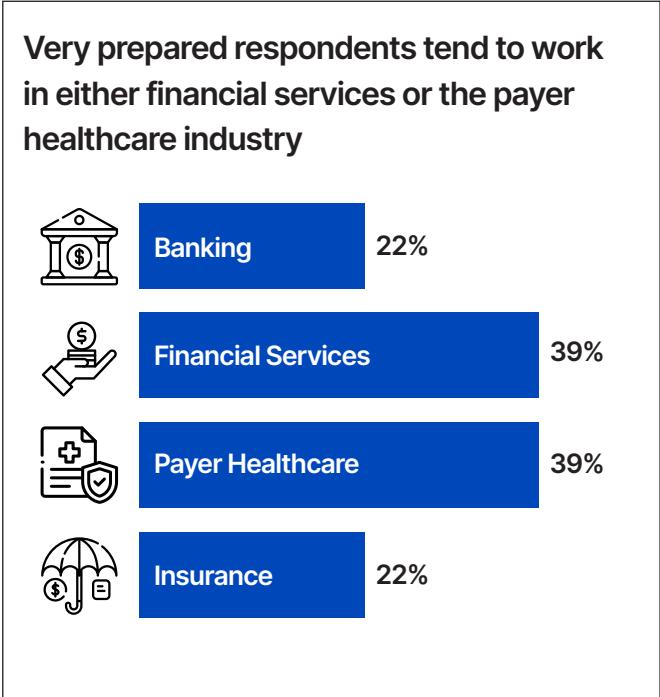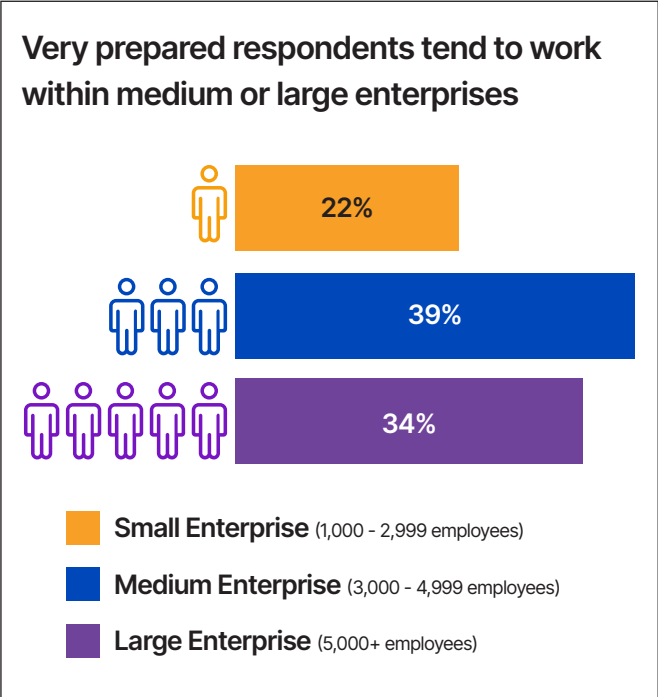The results of our analysis show that very few organizations are fully prepared to begin benefiting from LLMs.

In fact, **just 30% of senior tech professionals are very prepared to run LLM projects for enterprise solutions**, while 70% are only somewhat prepared or unprepared.

Very Prepared
30% of respondents

Somewhat Prepared or Unprepared
70% of respondents
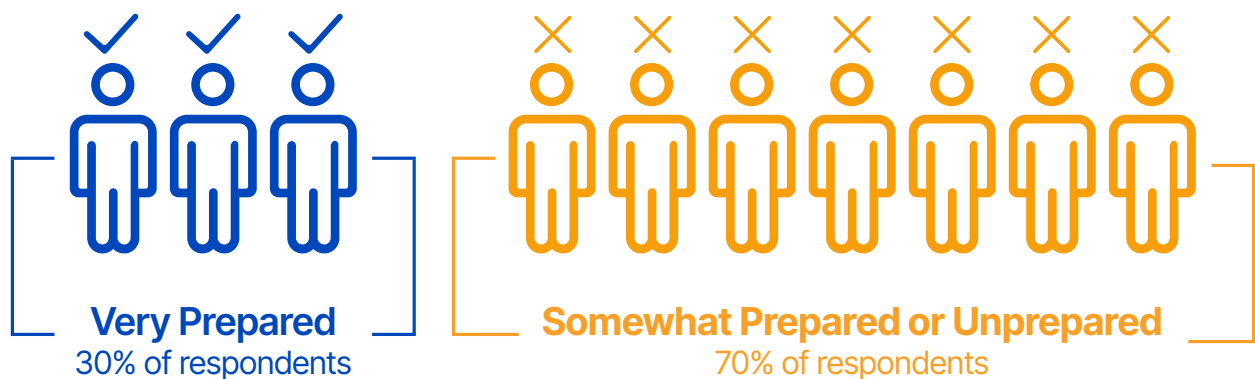
Notes:
[1]"2-5', '6-9' and '>10' responses combined.

Let's take a look at the characteristics of tech professionals who are very prepared to start benefiting from integrating LLMs into their business applications.

Very prepared respondents tend to work within medium or large enterprises

- 22%
- 39%
- 34%

■ **Small Enterprise** (1,000 - 2,999 employees)
■ **Medium Enterprise** (3,000 - 4,999 employees)
■ **Large Enterprise** (5,000+ employees)

Very prepared respondents tend to work in either financial services or the payer healthcare industry

Banking	22%
Financial Services	39%
Payer Healthcare	39%
Insurance	22%

IT and solutions professionals, including enterprise architects, are more likely to be very prepared than machine learning or AI specialists

Job role	% of very prepared respondents
IT and solutions professionals	**34%**
Machine Learning / AI specialists	**25%**

The stats show that senior tech professionals in medium / large enterprises, with IT and solutions roles and/or working in the financial services industry are most likely to have their ducks in a row when it comes to adopting LLMs.

On the other hand, small enterprises, those in machine learning / AI job roles and those in the insurance, payer healthcare or banking industries are falling behind.

However, the fact that relatively low percentages of senior tech professionals are confident that they are very prepared to integrate LLMs shows that everyone, no matter their company size, job role or industry, illustrates that there is foundational work to do if they wish to harness and scale the power of LLMs.

What is also clear is that something is holding tech professionals back from moving forward with LLMs, as we explore in the following section.

6

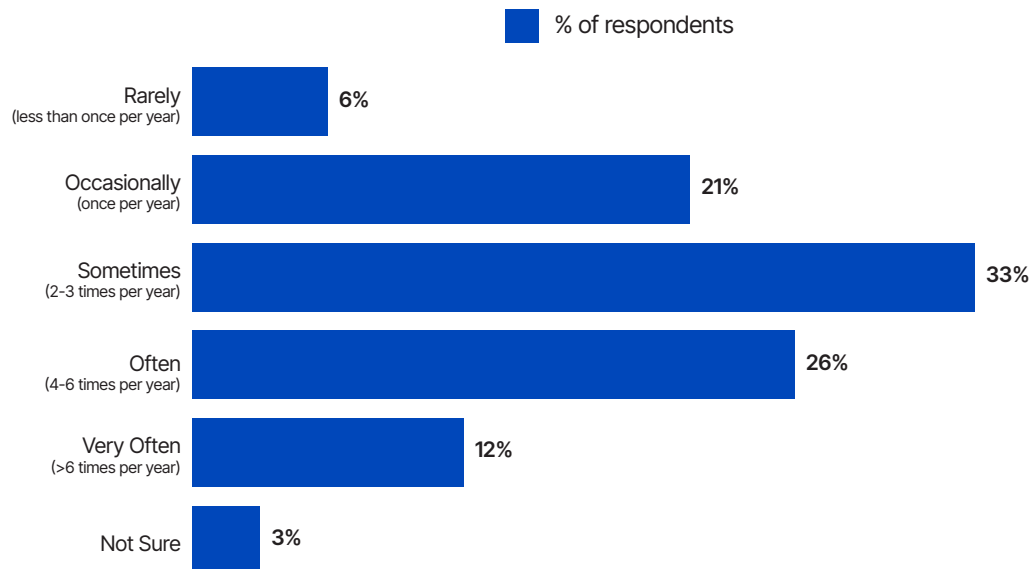# Why AI Projects Aren't Going Beyond Experimentation

The majority of senior tech professionals surveyed recognize that their organization would benefit from integrating LLMs.

However, in reality, only a few organizations plan to initiate many LLM projects in the coming years.

In fact, 60%[1] of senior tech professionals surveyed plan to initiate up to three LLM-integrated projects per year in the next two years with only 37%[2] expecting four or more projects per year.

Clearly, something is holding some tech teams back from operationalizing LLM projects.

Q. In the next two years, how frequently do you plan to initiate LLM projects for enterprise solutions?

■ % of respondents

Category	%
Rarely (less than once per year)	6%
Occasionally (once per year)	21%
Sometimes (2-3 times per year)	33%
Often (4-6 times per year)	26%
Very Often (>6 times per year)	12%
Not Sure	3%

The research shows that, in fact, almost all (95%)[3] senior tech professionals surveyed are encountering one or more challenges when it comes to their LLM projects, with just 5% saying that nothing is preventing them from going from experimentation to production.

According to the findings, tech teams are most likely to be troubled by security concerns (41%), while resource limitations (28%), difficulty proving ROI (27%) and a range of other challenges are halting progress.
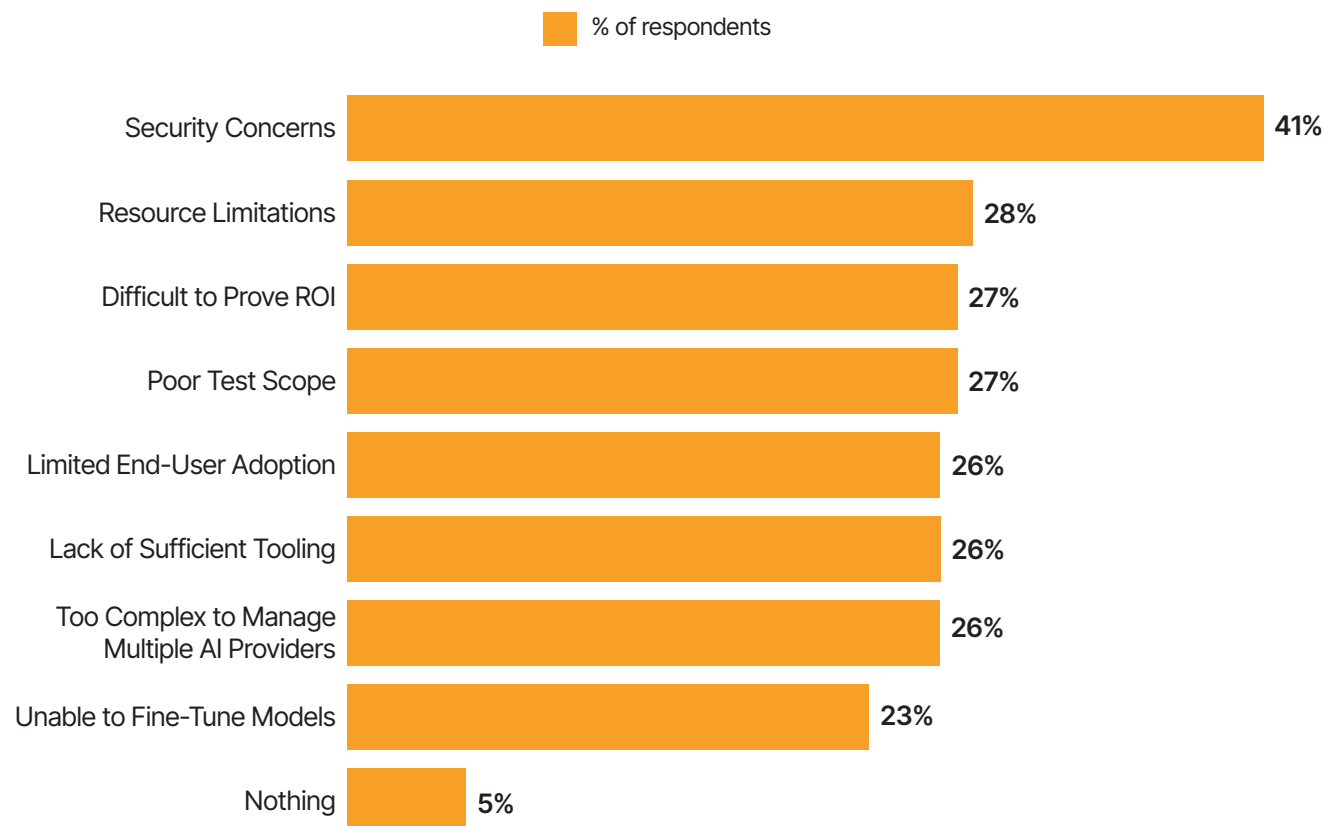
Notes:
[1]'Rarely (less than once a year), 'Occasionally (once a year)' and 'Sometimes (2-3 times a year)' responses combined.
[2]'Often (4-6 times a year)' and 'Very often (more than 6 times a year)' responses combined
[3]'Inverse of 'Nothing is preventing AI project from going from experiment to production'.

Q. What, if anything, are the top reasons preventing AI projects from going from experiment to production? (select up to three)

■ % of respondents

Security Concerns	41%
Resource Limitations	28%
Difficult to Prove ROI	27%
Poor Test Scope	27%
Limited End-User Adoption	26%
Lack of Sufficient Tooling	26%
Too Complex to Manage Multiple AI Providers	26%
Unable to Fine-Tune Models	23%
Nothing	5%

Further findings also indicate that other factors including interoperability, performance, compatibility with legacy systems and team skills are causing headaches when it comes to integrating LLM solutions with existing enterprise infrastructure.

Q. How challenging, if at all, do you find the following when it comes to integrating LLM solutions with existing enterprise infrastructure?

Compatibility with Legacy Systems
75% of respondents[1]

Security
73% of respondents[1]

Interoperability
71% of respondents[1]

Team Skills
71% of respondents[1]

Performance
70% of respondents[1]

Notes:
[1] 'Extremely challenging' and 'Somewhat challenging' responses combined.

Another factor that may be causing tech teams headaches as they attempt to operationalize LLMs is that often projects will involve more than one inference provider.

We found that just over two thirds (67%) of senior tech professionals surveyed think they will have multiple inference providers, which could be proving difficult to manage.

67% think they will have more than one inference provider

On the whole, the research reveals that most tech leaders are experiencing challenges on their journey to adopt LLMs.

However, organizations that are furthest ahead in their preparations appear less likely to let these challenges stand in the way of initiating LLM projects.

In fact, we'll see that by laying the foundations for integrating LLMs into their business applications, despite the challenges involved, very prepared tech teams are maximizing their potential to benefit from using this technology.

Maximizing the Benefits of LLMs

Our data suggests that by ensuring their organization is ready to hit the ground running with LLMs in the coming years, the most prepared tech teams may be more likely to reap various rewards than those that have so far stayed away from the technology.

This is indicated by the fact that senior tech professionals surveyed who are very prepared (63%) are much more likely to believe fine-tuned LLMs would bring significant value to their organization than those who are only somewhat prepared (35%) or unprepared (28%).
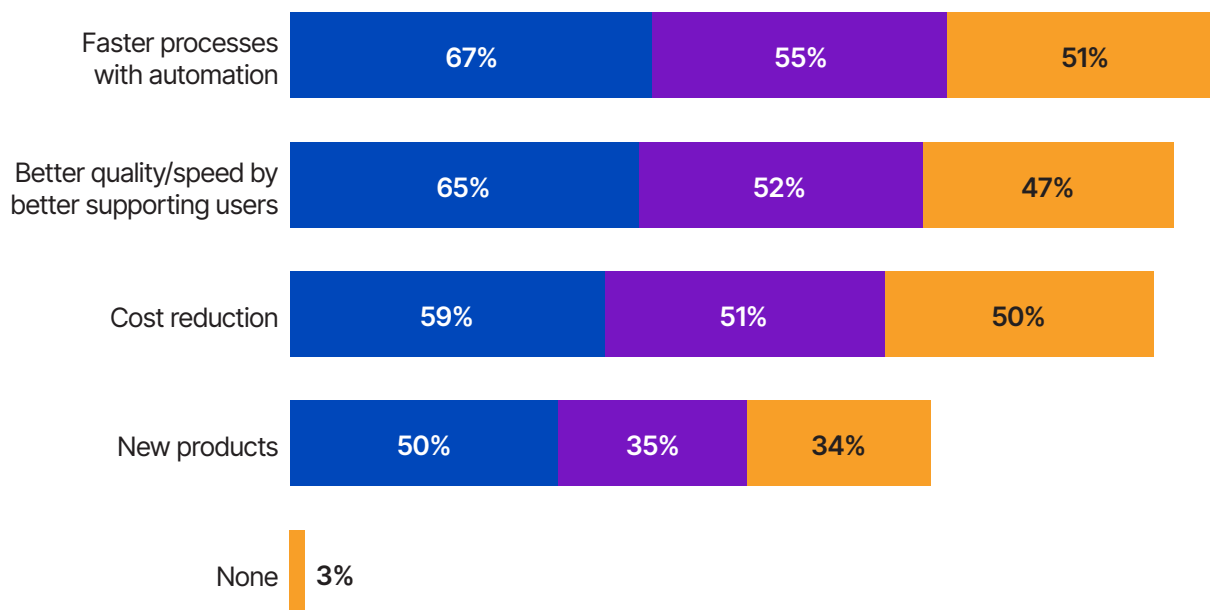
Fine-tuned LLMs bring value

9 out of 10 of senior tech professionals surveyed say fine-tuned LLMs would bring value[1] to their organization.

Q. If you had a 99% accurate LLM for a specific set of tasks, which benefit(s), if any, would you expect to see within the first 12 months?

■ % of very prepared respondents　　■ % of somewhat prepared respondents　　■ % of unprepared respondents

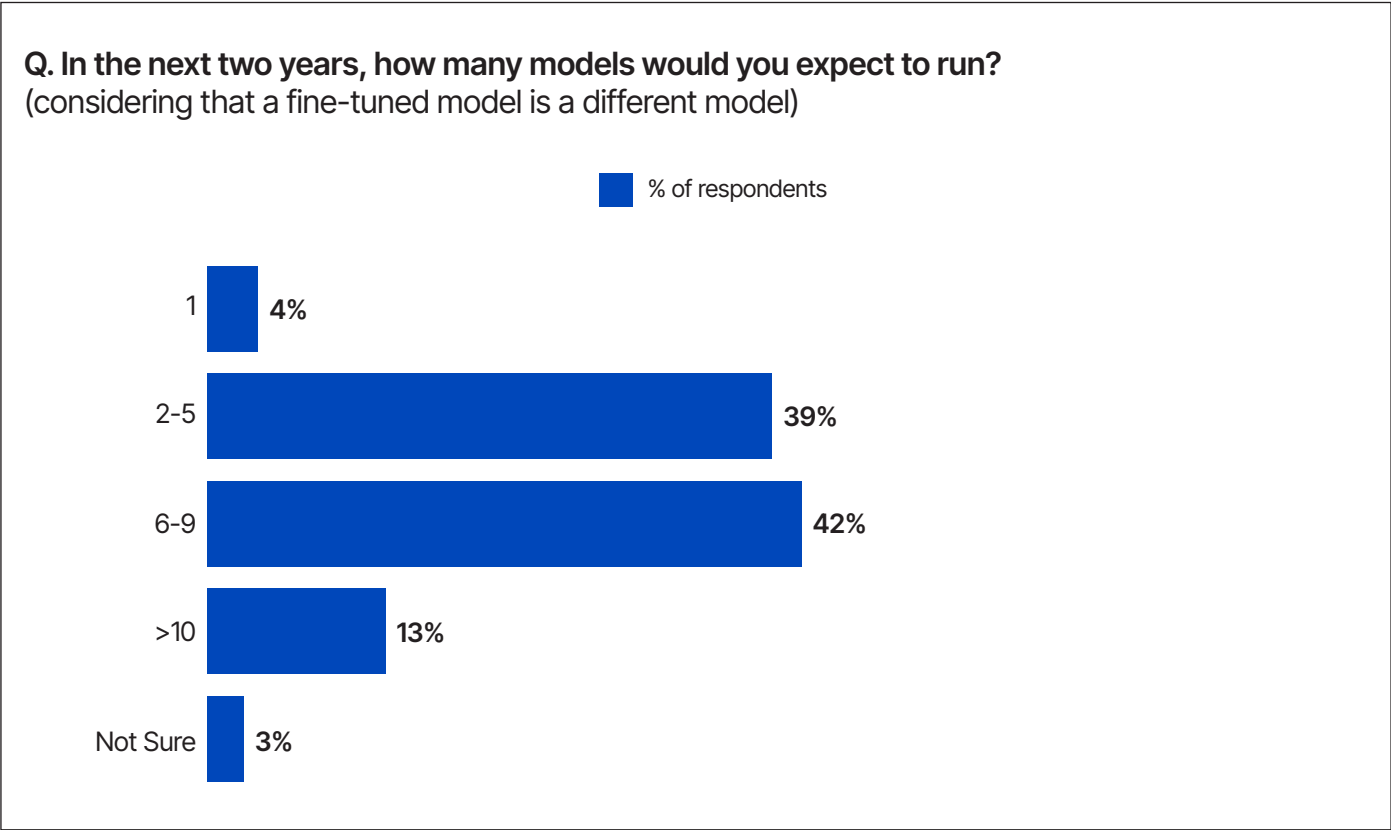|			
Faster processes with automation	67%	55%	51%
Better quality/speed by better supporting users	65%	52%	47%
Cost reduction	59%	51%	50%
New products	50%	35%	34%
None	3%		

These findings make a compelling case for doubling down on preparations for integrating LLMs into business applications.

Next, we explore what the most prepared organizations have already been doing to give themselves the best chance of success.

Preparing for Success in a Multi-Provider LLM World

The findings suggest that tech teams who are very prepared to integrate LLMs into their business applications understand that relying on a single GenAI model won't achieve the best results. In fact, the most prepared senior tech professionals expect to run more models than their less prepared counterparts.

Very prepared senior tech professionals expect to run nine models in the next two years[1] compared to six models by somewhat prepared respondents, and three models by unprepared respondents.

Q. In the next two years, how many models would you expect to run?
(considering that a fine-tuned model is a different model)

■ % of respondents

1	4%
2-5	39%
6-9	42%
>10	13%
Not Sure	3%

Notes:
[1] Mean: (No. of models excl. "Not sure")

The study also shows that very prepared tech professionals have been busy exploring various foundational models.

As shown by the chart below, the most prepared tech teams are more likely than those who are unprepared to have tested or are interested in trying out all the foundational models mentioned in our survey.

Q. Which foundational model(s), if any, have you tested or are you interested in?

█ % of very prepared respondents ██ % of unprepared respondents

Model	% of very prepared respondents	% of unprepared respondents
LLama	73%	59%
Titan	82%	63%
Cohere	79%	56%
Gemini	87%	68%
Claude	79%	63%
Mistral	73%	59%
OpenAI	96%	86%
Apple MM1	89%	68%
AI21	81%	64%
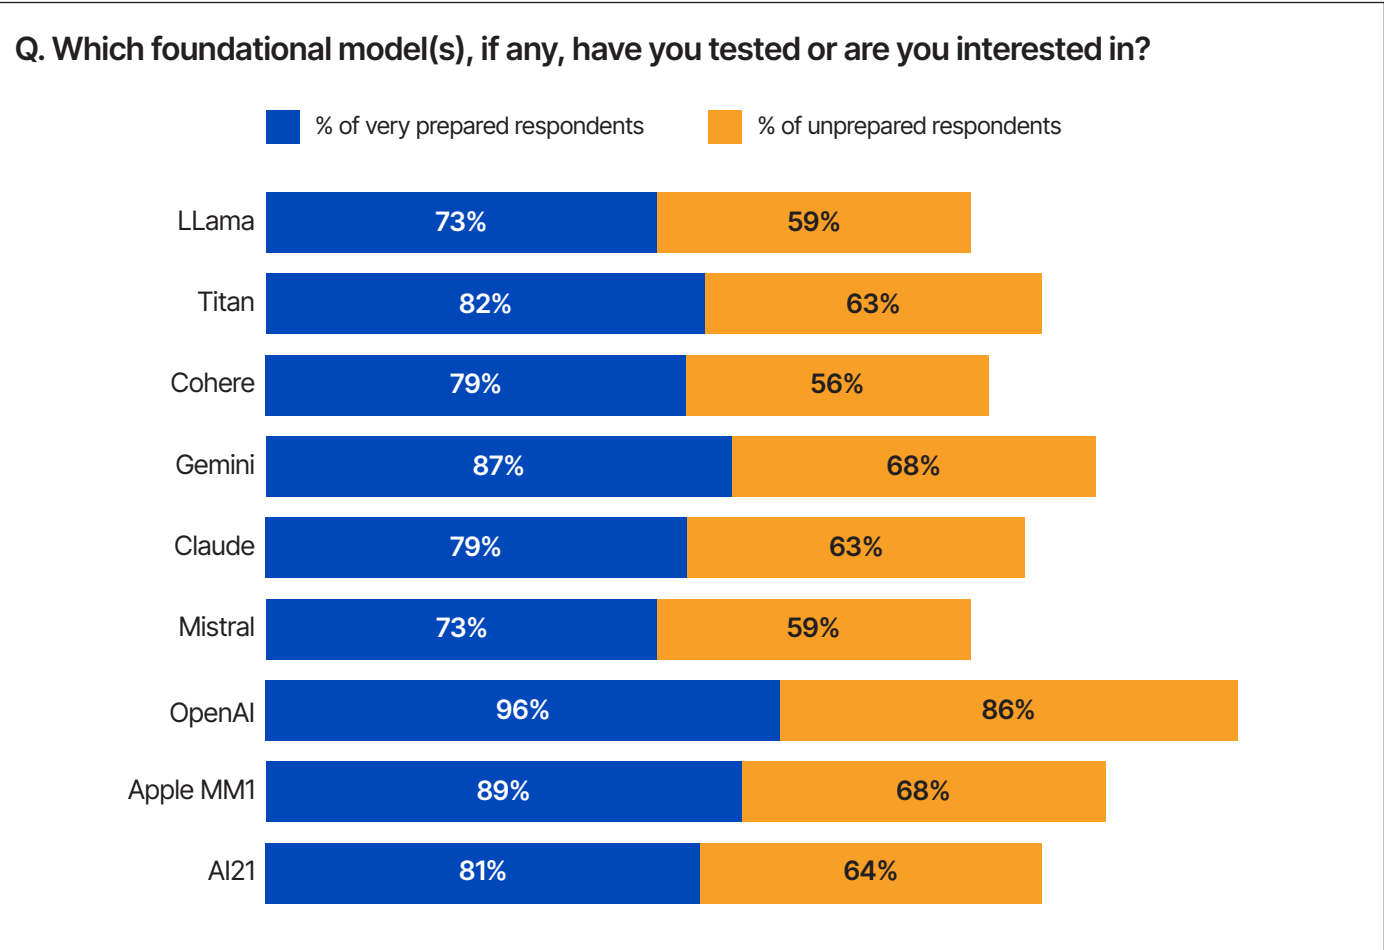
It's also worth noting that very prepared tech professionals are much more likely than those who are unprepared to look into lesser-known models such as Cohere or Apple MM1. On the other hand, unprepared respondents appear to be focusing on more familiar models like OpenAI and may be missing out on models that are better suited to their needs.

Overall, the findings suggest that **those who are ready to adopt LLMs don't expect to rely on just one foundational model**. Instead, they understand the importance of not being beholden to one particular model and finding a tool that's the best fit for the specific set of tasks it's expected to perform. This flexible mindset is further demonstrated by the fact that the **very prepared senior tech professionals surveyed are most likely to say that they expect they will have more than one inference provider** (80%).

Not only does this demonstrate the importance of experimenting with a range of providers as part of preparations for adopting LLMs, but it also highlights the need to choose a platform for building AI/LLM applications that integrate with multiple models and providers.

There is also evidence to suggest that senior tech professionals surveyed who are very prepared are doing their research when it comes to LLM platforms. In fact, this group is always more likely than their less prepared counterparts to think that various platform features are important, indicating greater levels of understanding.

In particular, very prepared tech leaders surveyed are most likely to rate virtualized LLMs (with load balancing, failover, multi-head execution) as important[1] (86%), while those who are unprepared are least likely to think these are important (57%).

It may be the case that those who are unprepared simply don't understand what virtualized LLMs are and the benefits they bring.

For example, virtualized LLMs can distribute tasks across multiple models to eliminate any one model as a single point of failure. Virtualized LLMs can also send tasks to multiple models in parallel as an approach to assess and select the best result, evaluate and gradually roll-out new models, or fine-tune lower cost models based on the results of better performing models.
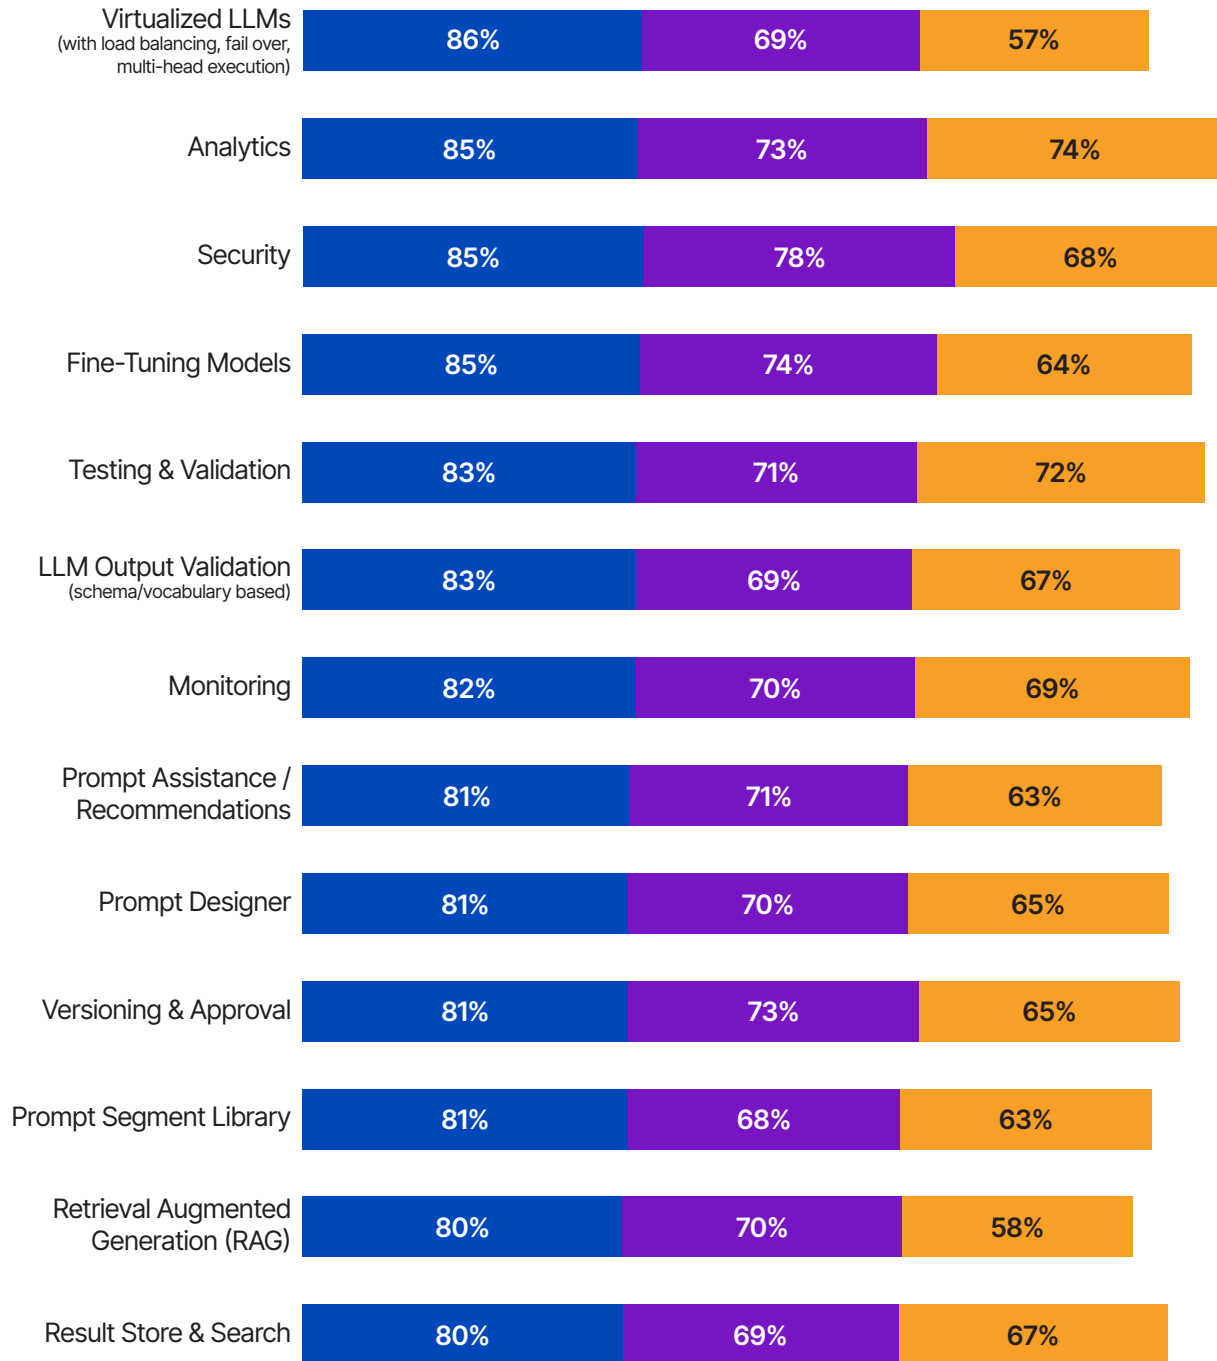
This is a great example of just how far ahead very prepared tech teams are when it comes to their understanding of LLMs.

Virtualized LLMs are important

86% of very prepared respondents rated Virtualized LLMs (with load balancing, fail over, and multi-head execution) as an important[1] feature

Notes:
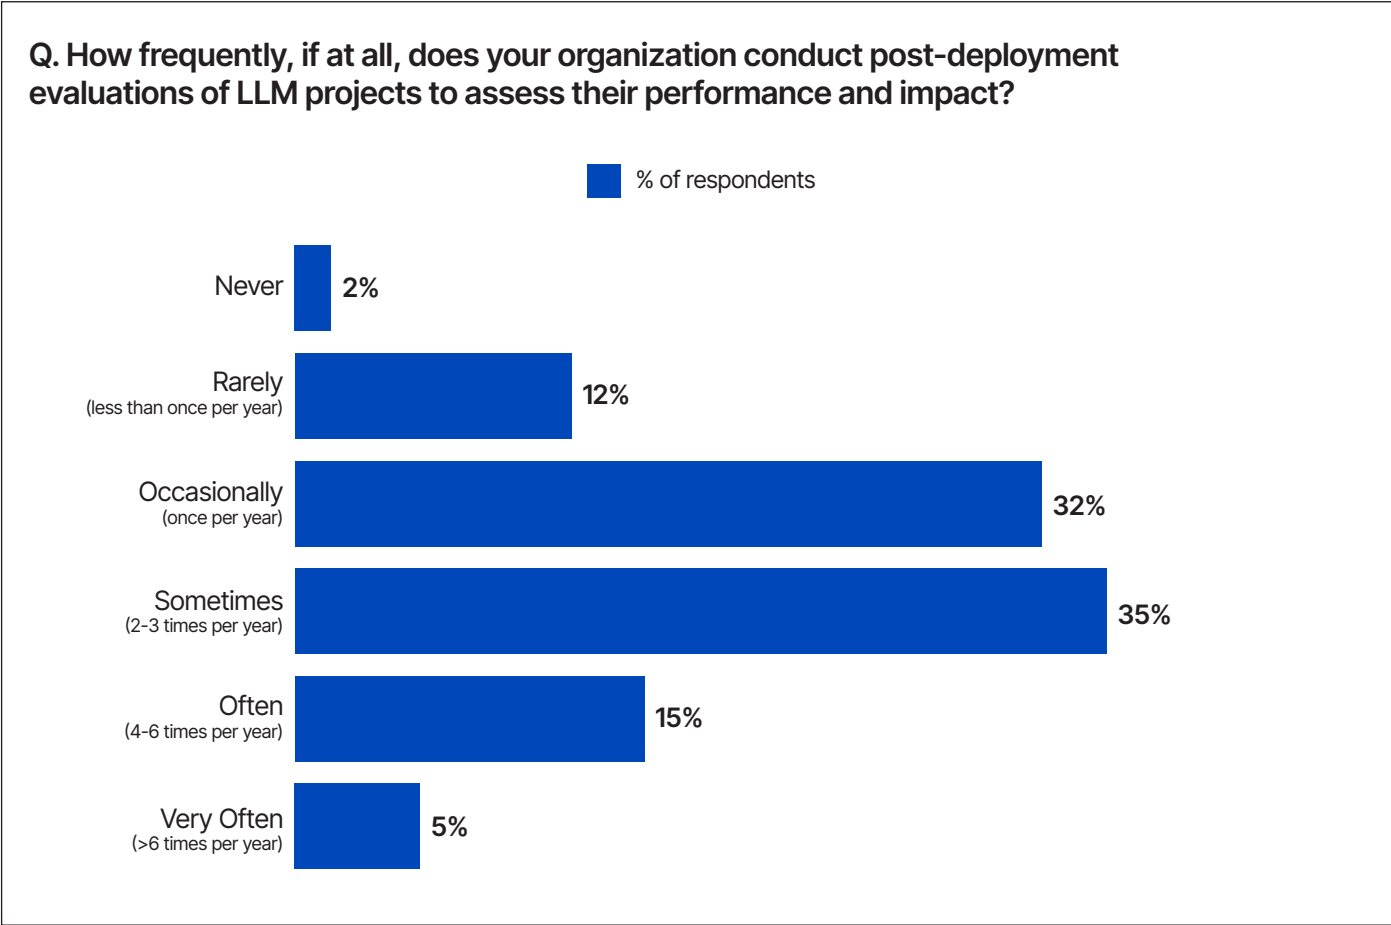[1] 'Very important' and 'Somewhat important' responses combined.

Q. How would you rate the importance of the following features of an LLM Operations Platform?

% of very prepared respondents who rate this feature as important (blue)
% of somewhat prepared respondents who rate this feature as important (purple)
% of unprepared respondents who rate this feature as important (orange)

Feature	Very prepared	Somewhat prepared	Unprepared
Virtualized LLMs (with load balancing, fail over, multi-head execution)	86%	69%	57%
Analytics	85%	73%	74%
Security	85%	78%	68%
Fine-Tuning Models	85%	74%	64%
Testing & Validation	83%	71%	72%
LLM Output Validation (schema/vocabulary based)	83%	69%	67%
Monitoring	82%	70%	69%
Prompt Assistance / Recommendations	81%	71%	63%
Prompt Designer	81%	70%	65%
Versioning & Approval	81%	73%	65%
Prompt Segment Library	81%	68%	63%
Retrieval Augmented Generation (RAG)	80%	70%	58%
Result Store & Search	80%	69%	67%

Finally, the most LLM-ready tech teams are preparing for projects in the coming years by dedicating more time to reviewing performance now.

The findings show that on average[1] those who are very prepared to adopt LLMs are already conducting post-deployment evaluations for projects four times per year, while those who are just somewhat prepared or unprepared only do this twice a year.

Q. How frequently, if at all, does your organization conduct post-deployment evaluations of LLM projects to assess their performance and impact?

■ % of respondents

Never	2%
Rarely (less than once per year)	12%
Occasionally (once per year)	32%
Sometimes (2-3 times per year)	35%
Often (4-6 times per year)	15%
Very Often (>6 times per year)	5%

Notes:
[1] Mean: (Times a year)

How are LLMs Being Deployed in the Enterprise?

By following the example of GenAI ready teams, and dedicating more time to researching, diversifying and reviewing their approach to integrating LLMs into their business applications, organizations will put themselves in a much better position to run successful LLM projects in future.

However, there's one thing that even the most prepared organizations haven't done yet. And that's to think outside the box.

According to the research, organizations that are planning to augment parts of their operation with LLMs in the next 12-18 months are most likely to say they plan to enhance their customer service and support (48%). This is especially likely to be the case among senior tech professionals surveyed that are very prepared to adopt LLMs (66%).

However, the use of AI is already well established in this area, perhaps suggesting a lack of imagination when it comes to the possible applications for LLMs.

Indeed, respondents were least likely to say they are planning to use LLMs for product development and engineering (prototyping, design optimization, virtual simulations for testing) (29%) in the next 12-18 months.

Q. Which of the following areas, if any, are you planning to augment with LLMs in the next 12-18 months?
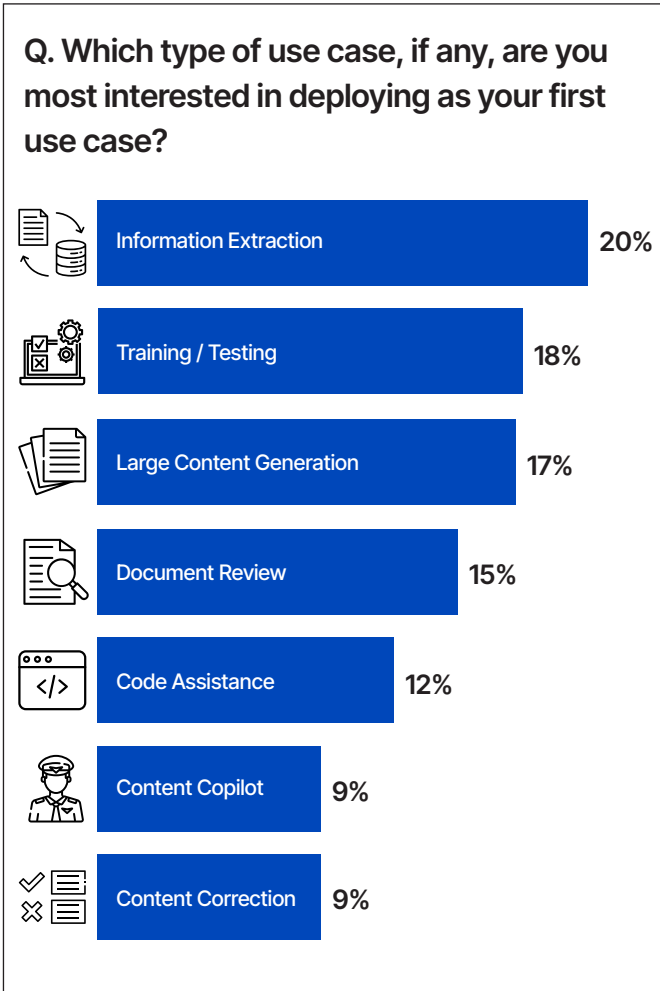
Area	%
Customer Service & Support	48%
Financial & Risk Management	46%
Supply Chain & Logistics	39%
Human Resources	37%
Marketing	33%
Legal & Compliance	32%
Product Development	29%

What's more, very few senior tech professionals surveyed seem keen to deploy LLMs in a variety of ways.

For example, just a fifth (20%) of respondents say they are most interested in deploying LLMs for information extraction as their first use case, despite this being the top response.

Meanwhile, even fewer senior tech professionals surveyed say they are most interested in using LLMs for training & testing (18%), large content generation (17%) and document review (15%) as their first use case.

Perhaps in time, as models mature and leveraging this disruptive technology becomes second nature to tech teams, we'll see more organizations forging ahead into new territory.

With the right platform to support them, the sky really is the limit.

Q. Which type of use case, if any, are you most interested in deploying as your first use case?

Use case	%
Information Extraction	20%
Training / Testing	18%
Large Content Generation	17%
Document Review	15%
Code Assistance	12%
Content Copilot	9%
Content Correction	9%

Key Takeaways

Enterprise teams expect to leverage multiple models and inference providers.

Working with multiple models and inference providers can add a layer of complexity and mean more time spent updating API endpoints. It may also limit experimentation and cause vendor lock-in, which could prevent teams from finding the best model for the task they wish to perform.

Very few organizations are fully prepared to leverage LLMs.

The data reveals that just 30% of senior tech professionals surveyed are very prepared to adopt LLMs, while 70% are only somewhat prepared or unprepared. So, you may not be as behind in your efforts as you think.

Preparing to integrate LLMs into enterprise solutions now will maximize your chances of experiencing a range of benefits.

We discovered that those who are most prepared are most likely to expect to achieve faster processes, better quality/speed, cost reduction and new products within the first 12 months of deploying a 99% accurate LLM for a specific set of tasks.

Tech professionals see the value of LLMs.

In fact, 90% of those we surveyed say fine-tuned LLMs would bring value[1] to their organization.

Tech professionals aren't thinking outside the box yet.

In fact, very few say they are considering deploying LLMs in new ways including as a code assistant, content copilot or content corrector. However, with time, experience and a little imagination, we're certain that we'll see more innovative use cases in future.

Notes:
[1] 'Significant value' and 'Some value' responses combined.

Conclusion

As companies increasingly incorporate Large Language Models (LLMs) into existing business processes and applications, Vertesia, an API-first platform, enables enterprise architects and traditional development teams to deliver significant business value with generative AI.

The platform simplifies the integration of diverse AI technologies, allowing businesses to select, switch, and manage LLMs and inference providers with unprecedented flexibility.

Vertesia is designed to accelerate the deployment of AI solutions, making it an essential tool for companies in sectors such as insurance, banking, retail, and technology, where leaders seek to improve business performance by automating and augmenting their content applications and processes.

Learn more at vertesiahq.com