

Unlocking enterprise AI: opportunities and strategies

A global study of 1,100 C-suite executives and
technologists, with interviews from 28 CIOs



Contents

3	About this report
4	Survey
5	Acknowledgments
6	Foreword by Andy Kofoid, president of global field operations, Databricks
7	Executive summary
10	Chapter one: AI's unlock moment
11	Executives, and especially practitioners, believe AI has proven value
12	Empowering everyone
14	The right decision, faster
16	A risk-graded approach
17	Test beds
19	Chapter two: Rewards and returns
19	Rational exuberance?
22	Data intelligence
23	Measuring tangible gains from AI
25	The virtue of patience
27	Chapter three: The infrastructure refit
28	Diseconomies of scale
30	Getting data right—and getting the right data
33	The burden of housekeeping
34	Practitioner perspectives
36	Chapter four: Find your edge
36	Mix and match
38	Choosing models
42	Chapter five: Guardrails and governance
44	Centres of excellence: 'forcing functions' for co-ordination
46	Calibrating humans and machines
50	Conclusion

About this report

This Economist Impact report, commissioned by Databricks, combines a global survey of 715 technical executives and 385 data and artificial intelligence (AI) technologists who work across the fields of data engineering, data science and enterprise architecture.

It also features insights from interviews with 28 C-suite executives from leading organisations across 11 industries: financial services, healthcare and life sciences, retail and consumer goods, the public sector, manufacturing, transport, energy, technology, professional

services, media and entertainment, and telecommunications. Digital native businesses across various industries are featured. Most interviewees are from multinational companies with headquarters spanning ten countries in the Americas, Asia-Pacific and Europe.

The report provides a comprehensive assessment of AI implementation from both the executives and practitioners' points of view, revealing catalysts and bottlenecks and identifying strategies for success.

Survey

This study's global survey, developed and conducted by Economist Impact and commissioned by Databricks, polled 1,100 respondents across the following backgrounds. It was fielded between July and August 2024.

The survey respondents work in large enterprises (annual revenue of US\$500m or greater) or public sector organisations, located in 19 countries across North America, Europe and Asia-Pacific (including ASEAN).

Respondents' job titles:

- Chief information officer
- Chief technology officer
- Chief data/analytics officer
- Chief data specialist
- Chief enterprise/data architect
- SVP/VP/head of IT/data/engineering
- Data scientists (or similar)
- Data engineers (or similar)
- Enterprise architects (or similar)

Respondents' industries:

- Financial services, banking and insurance
- Public sector
- Healthcare, pharmaceuticals and life sciences
- Retail and consumer goods
- Manufacturing (including automotive)
- Media and entertainment
- Energy, oil and gas
- Telecommunications

Respondents' locations:

Americas

- United States

Asia-Pacific

- Australia
- India
- Japan
- Malaysia
- Philippines
- Singapore
- South Korea
- Thailand

Europe

- Denmark
- Finland
- France
- Germany
- Italy
- Netherlands
- Norway
- Spain
- Sweden
- United Kingdom

Acknowledgments

We would like to thank the following executives for participating in interviews and sharing their insights:

- **Chalee Asavathiratham**, former chief digital banking officer, **Siam Commercial Bank**
- **Wassym Bensaid**, chief software officer, **Rivian**
- **Sanjay Bhakta**, chief product and technology officer, **Condé Nast**
- **Ian Botts**, chief technology officer, **Fanatics Betting & Gaming**
- **Bernd Bucher**, chief information officer, **Novartis**
- **Roman Bugaev**, chief technology officer, **Flo Health**
- **Ting Cai**, chief AI and data officer, **Rakuten Group**
- **Juan Jose Casado**, chief digital officer, **Repsol**
- **Kushal Chakrabarti**, chief data officer, **Opendoor**
- **Helen Choi**, chief digital and information officer, **CJ CheilJedang**
- **Carol Clements**, chief digital and technology officer, **JetBlue**
- **Ken Finnerty**, president of IT and data analytics, **UPS**
- **Jon Francis**, chief data and analytics officer, **General Motors**
- **Leonel Garciga**, chief information officer, **US Army**
- **Scott Hallworth**, chief data and analytics officer, **HP**
- **Andy Hill**, chief data officer, **Unilever**
- **Mohit Kapoor**, group chief technology officer, **Mahindra Group**
- **Jeff Martin**, senior vice president and chief data officer, **TD Bank Group**
- **Melissa Pint**, chief digital information officer, **Frontier**
- **Senthil Ramani**, global lead, data and AI, **Accenture**
- **Gereurd Roberts**, group managing director, Seven Digital, **Seven West Media**
- **Robbert Van Rutten**, chief information officer, **Shell**
- **Takaaki Sato**, senior executive vice president, chief technology officer, **NTT Docomo**
- **Amit Sharma**, chief technology officer, **Dream Sports**
- **Ryan Snyder**, senior vice president and chief information officer, **Thermo Fisher Scientific**
- **Sara Vaezy**, chief strategy and digital officer, **Providence**
- **Greg Ulrich**, chief AI and data officer, **Mastercard**
- **Darrin Vohs**, global chief information officer, **Molson Coors**

Foreword by Andy Kofoid, president of global field operations, Databricks

The present trajectory of artificial intelligence (AI)—investment, adoption, and results—makes abundantly clear that the technology is going to become an integral part of every business, irrespective of industry. Just as certain is the fact that success in this new landscape will require business leaders to do far more than adopt the latest and greatest AI model.

Winners in each industry will be those who take a holistic approach that encompasses data management, security, governance, culture, and domain-specific expertise. All of these areas are covered in the findings from this Economist Impact report, which surveyed 1,100 executives and technologists across 19 countries and eight industries and interviewed another 28 C-level executives. As a result, it is packed with insights to help technology leaders and data and AI professionals navigate their own AI journeys.

A key finding from this report is that companies overwhelmingly want to integrate generative AI models with their own data. That requires modernising their enterprise data platform—addressing the problem of fragmentation and advancing beyond the capabilities of a single large commercial model. The future also lies in AI systems and agents comprising multiple components, using custom models, each optimised for specific tasks. This approach allows for greater accuracy, cost, performance and security—crucial factors when dealing with enterprise-grade, AI-powered applications.

This report also finds that democratisation is a big driver for AI. Organisations want to make data more accessible by a wider range of employees and look to empower their people to build and refine AI models that reflect their day-to-day business needs. Imagine being able to ask your data a simple question in English or your native language, like "How is our European business doing on fiscal year-end targets?" and getting immediate, detailed answers that are certified as accurate, truly reflecting the nuances of your business.

Security and governance are also essential in the new AI landscape. As organisations rush to implement AI solutions, they grapple with how to protect sensitive data and comply with evolving regulations. Sometimes, they move too fast, but in this world, it's also often detrimental to move too slow. Putting robust security and governance processes in place is critically important to unlocking pathways for innovation and not becoming paralysed by an overabundance of caution, particularly as the regulatory environment around AI evolves.

All of the above means that the future lies not only with AI models trained on web data, but with true data intelligence—ie, systems that deeply understand and operate on an organisation's specific, often sensitive data sets. AI that understands the context of your unique business. This approach—agent systems leveraging specialised models—allows companies to build solutions tailored to their particular needs and domains, leveraging their proprietary data as a powerful competitive advantage.

The AI revolution is just beginning, and one thing is certain: the organisations that effectively harness their data and build AI systems tailored to their unique needs will be the leaders of tomorrow. At Databricks, we're proud to support research that helps illuminate the opportunities and strategies that will define the AI era.

Andy Kofoid

President of global field operations, Databricks



Executive summary

A heady US\$1trn in capital expenditure is expected in the coming years to deliver the data centres, chips, energy and infrastructure to support the artificial intelligence (AI) revolution.¹ While this figure is largely driven by the dawn of generative AI (GenAI), the truth is that companies have been developing their AI capabilities for years. Breakthroughs like deep learning and neural networks have allowed specialists in fields from biotech to finance to crunch vast datasets to uncover patterns and deliver actionable intelligence.

The difference now is democratisation and scale. Because GenAI provides an intuitive, natural language interface, the benefits of AI have become accessible to every practitioner. To transition successfully from pilots to enterprise-wide deployment requires a robust infrastructure that can handle the data and computational requirements of AI, a workforce strategy that appropriately calibrates humans versus machines, and an appropriate return on investment (ROI) strategy. Last, but not

least, achieving AI excellence necessitates strong governance and careful design of human machine interaction. Decisions and outputs must comply with extensive existing and emerging legislation in areas like data privacy, security and consumer protection.

This Economist Impact report, commissioned by Databricks, combines a global survey of 715 technical executives and 385 data and AI technologists who work across the fields of data engineering, data science and enterprise architecture. It also features insights from interviews with 28 C-suite executives from leading organisations across ten countries, who represent 11 industries.

Blending survey trends with qualitative insights, this report provides a comprehensive assessment of AI implementation from both executives and practitioners' points of view, revealing catalysts and bottlenecks and identifying strategies for success.

1 Goldman Sachs, "Gen AI: too much spend, too little benefit?", 2024, <https://www.goldmansachs.com/insights/top-of-mind/gen-ai-too-much-spend-too-little-benefit>

Executives and practitioners alike believe in the power of AI, but think investment is falling short. From accelerating drug development to extending credit to the financially excluded, companies are finding a diverse array of applications for AI and GenAI. Our survey found that 85% of organisations are actively using GenAI in at least one business function, reaching 97% of companies with revenue over US\$10bn. IT teams are the most avid adopters and the legal function is the most reticent. Internal projects are preferred but, by 2027, 99% of executives expect GenAI adoption across internal and external use cases. Seven in ten see AI as crucial to their long-term strategic goals, and only 18% say it is overhyped. Despite the momentum, only one in five believe that their current investment across technical and non-technical domains is sufficient.

AI's short-term benefits include productivity and efficiency, but leaders will use AI to unlock value, revenue and business model innovation in the long term. So far, productivity is one of the most commonly reported impacts of AI. In line with this, functions with high automation potential have been heavily testing GenAI, such as IT (91%), marketing (85%), sales and customer service (83%), and operations (80%), while 82% of data scientists report using AI for coding. Interviewees also report a range of benefits, from cost reduction to improved employee experience and talent attraction. However, executives do take a 'long' view of the GenAI paradigm shift, noting that strategic

considerations like business model innovation, market positioning, and environment, social and governance standards have been the most important elements when evaluating business cases for AI. Revenue growth has been the least effective metric for justifying investment to date, with experts arguing that AI returns will take time to accrue given the need for experimentation, iteration and digital infrastructure overhauls. Over time, however, the ability to realise revenue will set leaders apart, and financial value will need to be more definitively quantified.


Enterprise-wide AI requires an infrastructure refit—most companies do not feel confident in their current architecture. Many enterprises are operating with the technological equivalent of Victorian-era plumbing. Only 22% of organisations say their current architecture can support AI workloads without modifications, and 48% of data engineers spend most of their time resolving data source connections. The problem could worsen as democratisation leads to a proliferation of AI pilots, leading to congestion, complexity, and opacity in data and infrastructure. The prize is worth the work; some of the companies interviewed for this programme reported that they could be truly creative in finding use cases and achieving returns only once they had secured their data foundation. Worries about data security, from silos to fragmented systems, are holding some companies back from more ambitious experimentation.

Two-thirds of organisations see significant potential in integrating GenAI models with their own data. As AI models become commoditised, the best performers will mix and match models and tools with their unique assets: proprietary data and know-how. This requires finding the right blend of tools, such as open- and closed-source models, and gaining visibility and mastery over internal datasets. Organisations are exploring a variety of models, and nearly seven out of ten are experimenting with or have fully deployed open-source GenAI, with 96% believing they will do so by 2027. Experts warn that off-the-shelf models are inferior in areas like domain-specific lexicon, and may provide less control and security. However, off-the-shelf models, unconnected to enterprise data, are still common: almost half (45%) of data scientists are using large language models (LLMs) without retrieval augmented generation (RAG).

Few enterprises have fully ‘productionised’ GenAI due to cost, skills and governance challenges. Only 37% of executives believe their GenAI applications are production-ready, a figure that falls to just 29% among practitioners. Data scientists cite key constraints including cost (41%), skills (40%), quality (37%) and governance (33%). Only one in six respondents believe that their organisation can secure enough AI talent, half of data engineers say governance takes up more time than anything else, and more than half of enterprise architects (53%) cite challenges with data privacy and security breaches as the biggest risk of AI expansion. Executives agree that in the years ahead, winners will be those that graduate from experiments to production, scaling and monetisation.

The inevitability of mass AI adoption underscores the critical need to carefully calibrate human and machine intelligence to enable democratised AI ecosystems.

Leaders say AI can augment rather than replace employees. Interviewees cite unified systems and ‘centres of excellence’ as critical to balancing governance and enablement, pointing to an irreplaceable role for experience and judgement in overseeing AI outputs. Two-thirds of organisations say they are still experimenting to find the right balance between humans and AI. Many are also investing in self-service tools and AI assistants to empower all employees to become data scientists. Enterprise architects see GenAI as a democratising force, with nearly 60% predicting that within three years natural language will be the primary—if not only—way that non-technical staff interact with complex datasets.



Chapter one:

AI's unlock moment

Organisations are racing to harness the power of AI, with generative models elevating capabilities across industries and functions. From enhancing drug discovery to powering intelligent assistants, a range of models and techniques are boosting productivity, guiding strategic decisions, and opening new frontiers of innovation. As companies move from experimentation and implementation, the future belongs to those who can successfully scale AI solutions.

Two years since the public launch of ChatGPT, every organisation is searching for algorithmic advantage. Companies are harnessing the powers of so-called traditional AI—which can recognise patterns, process vast and varied data sets, and make predictions based on past data—with the ability of GenAI to reason and converse across contexts, creating novel outputs that mimic human ability from natural language prompts.²

The GenAI breakthrough took AI from narrow tasks performed by expert teams to a capability available to nearly everyone. “Gen AI has exploded because of its approachability; the average person can use these tools,” says Darrin Vohs, global chief information officer at Molson Coors, a Canadian beverage company.

Many companies were already experimenting with AI long before ChatGPT was released, especially larger enterprises in more technically advanced fields. Novartis, a Swiss pharmaceutical giant, had been using machine learning (ML) for drug discovery and predictive analytics to improve clinical trials.^{3,4} Flo Health, a global consumer-facing women’s health app used by over 380 million people, is based on an algorithm using traditional ML techniques. “Advancements in LLMs have been really exciting for us,” says Andy Hill, chief data officer at Unilever, a consumer goods firm. “But it’s still the combination of that and the more traditional supervised learning techniques that together will provide the value for a company like us.”

2 MIT News | Massachusetts Institute of Technology, “Explained: Generative AI”, 2023, <https://news.mit.edu/2023/explained-generative-ai-1109>

3 Novartis, “The art of drug design in a technological age”, 2021, <https://www.novartis.com/stories/art-drug-design-technological-age>

4 Novartis, “Novartis’ commitment to the ethical and responsible use of Artificial Intelligence (AI) Systems”, 2020, https://www.novartis.com/sites/novartis_com/files/novartis-responsible-use-of-ai-systems.pdf

Generative models elevate AI capabilities due to their broad applications. Even those with ample AI experience have added a new layer of capability. Thermo Fisher Scientific, a global life sciences company, has been incorporating AI into its instruments for nearly a decade. GenAI builds on this strong AI/ML foundation, according to Ryan Snyder, the company's chief information officer, enabling deeper interrogation of data and bolder experimentation.

Executives, and especially practitioners, believe AI has proven value

While AI has been through its share of winters, our survey finds respondents are overwhelmingly optimistic, with 73% saying GenAI is critical to their long-term strategic goals. Only 18% say AI is overhyped, although executives are slightly more cautious (20% lean towards there being too much hype) than data and AI technologists (just 13%). In the years ahead, winners and losers will be defined by who is able to graduate from experiments to implementation at scale.

Amid this cautious optimism, the rapid adoption of GenAI is hard to ignore. An impressive 85% of organisations in our survey are testing and using it today, with adoption rates climbing as company size increases—among those with revenue over US\$10bn, 97% are testing and using it. This trend suggests that larger organisations may cement their competitive advantages through AI, while those lagging in investment may miss out on the next major technology shift. "There's always some truth underneath the hype. When the dot-com boom happened, lots of companies failed, but a lot of real companies with valuable products came out of it as well, and I believe the same is going to happen in the GenAI space," predicts Mr Vohs at Molson Coors. "GenAI is going to be a commodity embedded in everything, similarly to the internet, and the value will be in the business problems it helps solve."

"GenAI is going to be a commodity embedded in everything, similarly to the internet, and the value will be in the business problems it helps solve."

Darrin Vohs, global chief information officer, Molson Coors



Empowering everyone

To date, GenAI usage is most advanced among IT teams, our survey found, with 62% having fully deployed it. Marketing and sales, including customer service, are leading adoption ahead of other business teams (see figure 1). Companies interviewed for this study report quantifiable productivity gains among tech teams in particular.

Flo Health has seen a 24% boost in data engineering efficiency, with AI assistants reducing mistakes and accelerating document generation. Repsol, a Spanish multinational energy company, ran a GenAI experiment with 200 coders and more than 200,000 lines of code across 20 programming languages, finding improvements in productivity of up to 30%, with an average of 7%, depending on the technology. The quality of output is high too. Japanese technology company Rakuten has also reported that a significant number of coding outputs from its internal AI tool were accepted without modification.

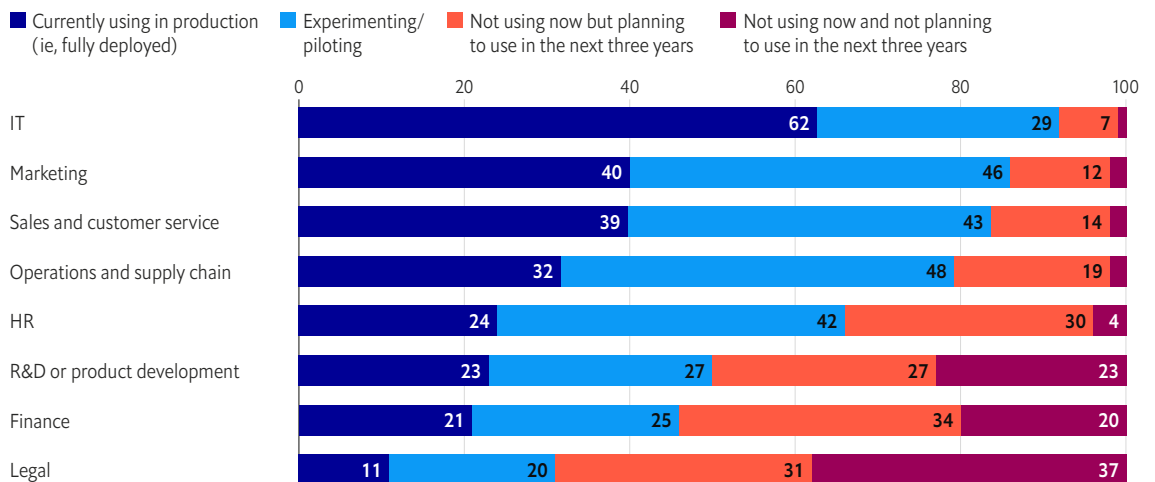
While those in tech roles like programming may find that GenAI is a natural fit, the idea of an assistant to everyone has increasing appeal.

‘Agentic AI’—artificial agents with a natural language interface that can plan and execute sequences of actions on behalf of a user—will become an increasingly common feature of daily life, surpassing text-to-speech applications like Alexa and Siri in scope and smarts. They could be holiday bookers, counsellors and creative sparring partners, or negotiate legal disputes and take on unscrupulous corporate practices.⁵ Our global survey finds that more than 97% of enterprise architects predict that non-technical staff will use natural language programmes to interact with complex datasets within the next three years. Nearly six in ten (58%) say that natural language will be either the primary or only way they do this.

The majority of business functions in our survey have either fully deployed or are actively experimenting with GenAI (see figure 1). The most hesitant in-house functions are legal and finance, the only divisions in our survey with less than half of respondents in full deployment or experimentation. That caution may be warranted given regulatory constraints and high stakes, although plenty of law firms are trying to embrace AI tools for use cases like drafting and reviewing contracts.

Figure 1: Stage of GenAI adoption by business function

Percentage of executives



Source: Economist Impact

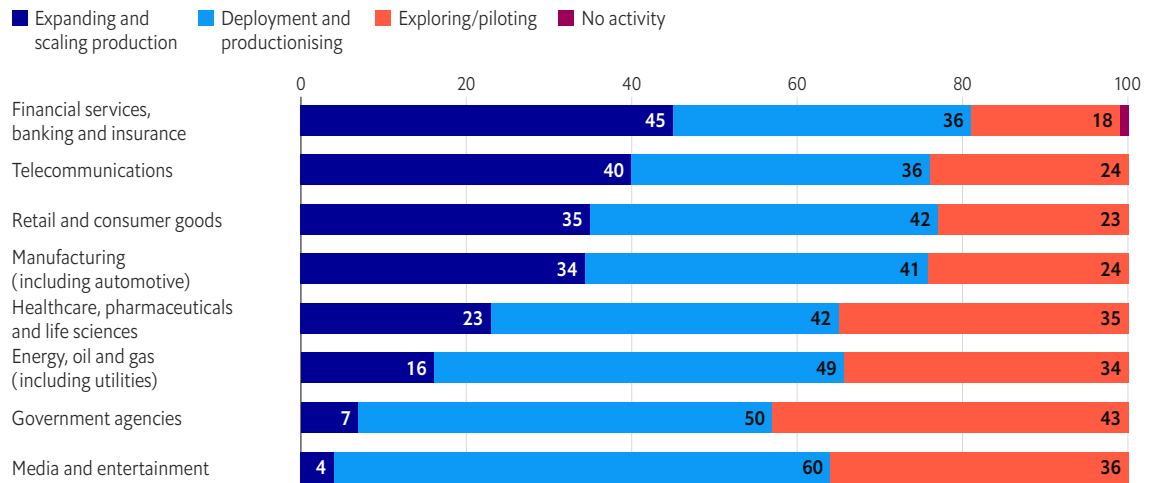
5 For example: <https://donotpay.com/>

In contrast to the trailing status of individual finance teams, financial services leads all other industries in GenAI implementation. For these firms, 45% are already expanding and scaling production of use cases and another 36% are in active deployment (see figure 2). Early adopters are also common across the telecoms, retail and

manufacturing sectors, suggesting significant potential for growth as they fine-tune their approaches. Finally, media and entertainment and the public sector currently lag in terms of scaled-up use cases, but the majority of organisations in those sectors are still beginning to actively deploy GenAI in production.

Figure 2: Stage of GenAI implementation by industry

Percentage of executives



Source: Economist Impact

In most countries, about 70% of organisations have moved past the piloting stage for GenAI (see figure 3). However, progress is much more varied when it comes to the scaling up stage.

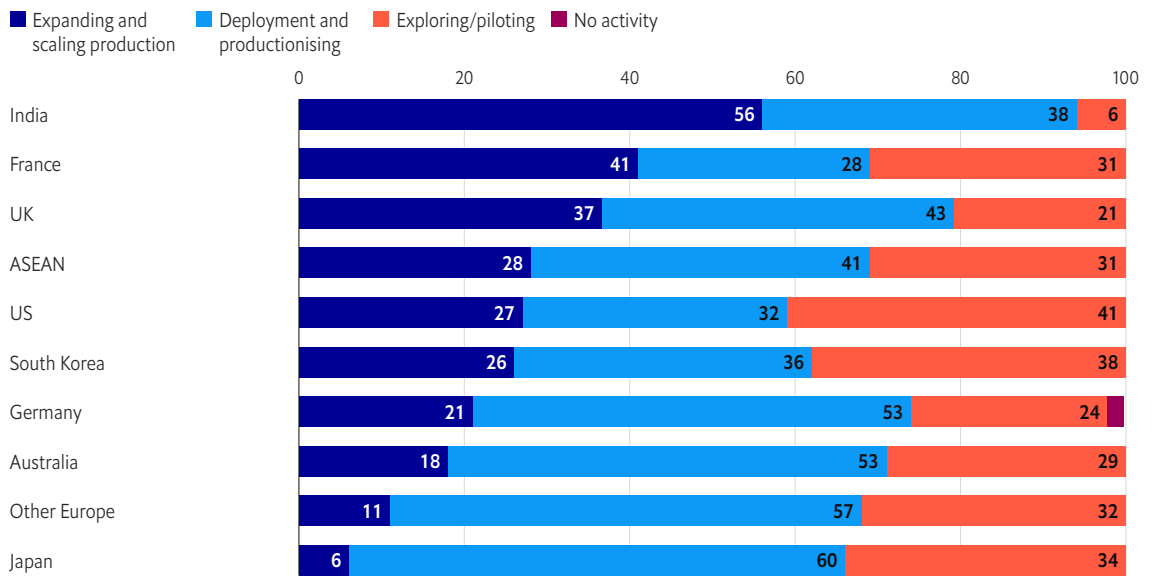
In India, for instance, most companies (56%) are scaling up GenAI, relying on public and open-source models much more than organisations in other countries.

Meanwhile, in Japan just 6% are scaling up, which may reflect the country’s preference for more complex, customised models. Half of Japanese executives said they were training custom models based on proprietary data, a rate far beyond any other country.

Similarly, for the US, which has the highest share of organisations still in the piloting stage, this perceived lag may not actually represent a weakness. Our data suggest that many of these experimenters may in fact be ‘GenAI leaders’—that is, they are also the most confident about the readiness of their GenAI applications. This is also true of firms in Europe to a lesser extent. These findings warrant further investigation, but it may be the case that leading companies in the US are more willing to take risks and explore more innovative use cases rather than rushing to scale up basic ‘tried and tested’ use cases.

Figure 3: Stage of GenAI implementation by country/region

Percentage of executives



Source: Economist Impact

ASEAN includes Malaysia, Philippines, Singapore and Thailand

Other Europe includes Denmark, Finland, Italy, Netherlands, Norway, Spain and Sweden

The right decision, faster

Productivity is an understandable initial focus because it's measurable—and it matters. But executives are using both classical AI/ML and GenAI to guide strategic decision-making in a volatile operating environment. Our survey found risk management was a dominant strategic area in which AI as a whole is informing decisions at companies of all sizes, reported by 68%, followed by market entry (66%), and strategic sales and partnerships (66%).

By contrast, larger companies are more likely to use AI for strategic decisions around product development (67%), hiring and workforce planning (67%), and investment and capital allocation (61%) compared with mid-size companies (55%, 53% and 48%, respectively). The value of AI for uncovering hidden insights is clear for many large enterprises, such as Mahindra

Group, an Indian multinational conglomerate. "We have many listed companies," says Mohit Kapoor, the group chief technology officer. "We ensure data is always identifiable so that it can be segregated, but we can also anonymise that data to drive insights from across our companies."

For firms operating in an increasingly unpredictable world, AI provides armour in the form of resilience and informs better decisions. UPS, a multinational shipping company, uses GenAI to help businesses across its 200 countries and territories on issues like understanding customs codes for products.⁶ "When we infuse decision science into software, we remove some of the noise and inconsistency with decision-making throughout a large worldwide company like ours," says Ken Finnerty, president of IT and data analytics at UPS. "We ensure consistency and raise quality. Delivering fast decisions, but more importantly the right decisions."

6 Mastercard, "Mastercard accelerates card fraud detection with generative AI technology", 2024, <https://www.mastercard.com/news/press/2024/may/mastercard-accelerates-card-fraud-detection-with-generative-ai-technology/>

“We’re scanning a trillion data points globally to predict whether a transaction is genuine or not, and we continue to improve the capabilities as we add new technologies. You can never determine that a transaction is fraudulent with 100% certainty, but we’re trying to give banks the confidence to make an informed decision.”

Greg Ulrich, chief AI and data officer, Mastercard

JetBlue, a US airline, has developed digital twins—ML-powered replicas of physical systems—to enhance predictive capabilities in an industry rife with unexpected disruptions from weather events to technical issues. “In the airline business, the number of unexpected things that can come your way in a given day is almost limitless, so our digital twin is becoming an incredible source of value for anticipating events and preparing solutions,” says Carol Clements, JetBlue’s chief digital and technology officer. This has improved operational efficiency and minimised disruption costs.

Mastercard uses GenAI to enhance and speed up card fraud detection. “We’re scanning a trillion data points globally to predict whether a transaction is genuine or not, and we continue to improve the capabilities as we add new technologies,” says Greg Ulrich, the company’s chief AI and data officer. “You can never determine that a transaction is fraudulent with 100% certainty, but we’re trying to give banks the confidence to make an informed decision.”

General Motors (GM) uses ML to analyse diagnostic data to predict issues before they become serious: the software can anticipate if a vehicle battery on the production line will have problems, for instance, so it can be removed

and fixed before being installed in a car. “It helps from a customer loyalty perspective as well, by averting a customer having a bad experience with their first EV [electric vehicle], which could erode trust in the brand,” says Jon Francis, the automaker’s chief data and analytics officer.

AI can also guide decisions that let companies expand their customer base without taking undue risk. Mahindra Group uses AI to analyse publicly available and proprietary data to make lending decisions, creating credit scores that could promote financial inclusion. Siam Commercial Bank, a Thailand-based bank, used ML to reduce credit underwriting costs by 20%, says Chalee Asavathiratham, the former chief digital banking officer. “Applying for a loan becomes unlocking, if you will, a pre-calculated credit number. So the loan process becomes much more streamlined, much easier.”

Unilever is embedding AI-driven intelligence into its decision-making process. Through the Horizon3 Lab, it is focusing on identifying and testing new AI concepts, designs and projects.⁷ Its Sky programme is leveraging advanced analytics to optimise its product portfolio to make faster, smarter decisions about what it puts on the shelves, explains Mr Hill. “I think we’ve made really significant progress towards being a data intelligent company, and there is still a huge opportunity ahead.”

“I think we’ve made really significant progress towards being a data intelligent company, and there is still a huge opportunity ahead.”

Andy Hill, chief data officer, Unilever

⁷ Unilever, “Unilever Launches Global AI Lab in Toronto”, 2023, <https://www.unilever.ca/news/press-releases/2023/unilever-launches-global-ai-lab-in-toronto/>

A risk-graded approach

Whether in productivity or in prototyping, companies are picking their pilots carefully. Sara Vaezy, chief strategy and digital officer for Providence, a healthcare organisation based in the US, explains its approach of being “very methodical and very intentional” to avoid the mindset of “spray and pray”. For its approach, she explains, “we set up a governance structure, we set up a technical environment where we could do these experiments very methodically and intentionally and instrument [them] properly, so we could understand what was happening.”

Many companies prefer internal-facing applications before releasing products to customers. TD Bank Group, a Canadian multinational financial institution, for instance, has deployed a GenAI chatbot for its call centre staff but has kept it internal. Firms are understandably cautious about releasing chatbots directly to customers. While they might offer speed and efficiency, unpredictability can lead to offensive content and misinformation, damaging both consumer trust and corporate reputations.

“We set up a governance structure, we set up a technical environment where we could do these experiments very methodically and intentionally and instrument [them] properly, so we could understand what was happening.”

Sara Vaezy, chief strategy and digital officer, Providence

Consumer perceptions are one consideration for firms looking to external use cases, to avoid convenience sliding into lack of control. Users may grow wary when AI feels invasive or manipulative, raising worries about privacy risks,⁸ or when it offers erratic recommendations, like telling users to eat rocks and glue cheese to pizza.⁹ Rogue bots have even been gamed to promote competitors.¹⁰ “With information that a user might act on, such as ‘what could I learn about a symptom I’m experiencing?’, there is undoubtedly a need for more caution,” says Roman Bugaev, chief technology officer at Flo Health.

The US Department of Defense (DoD) draws a line between administrative and battlefield applications, with far higher stakes associated with decision support in warfare. “It’s one thing using an LLM on the administrative side, to process DoD regulations to make sure your documents meet them. But decision support in warfare, like identifying a bad guy to shoot, is higher stakes,” says Leonel Garciga, chief information officer for the US Army. Decision support will still remain with the respective commander, and humans will always be in the loop as far as AI is involved.

Internal experiments are a safe testing ground. “If you introduce products without high reliability, that may erode customer trust,” argues Wassym Bensaid, chief software officer at Rivian, an EV manufacturer. “Launching a product requires a steep maturation climb from demonstration software. We only introduce products once we’ve addressed all the safeguards from a quality and reliability standpoint to ensure the technology will be a trusted assistant for our customers.”

8 BBC, “Snapchat’s My AI chatbot: The privacy concerns”, 2023, <https://www.bbc.co.uk/news/technology-67027282>

9 BBC, “Glue pizza and eat rocks: Google AI search errors go viral”, 2024, <https://www.bbc.co.uk/news/articles/cd11gzejgz4o>

10 Carscoops, “Chevy Dealers’ ChatGPT Bots Recommend Teslas, BMWs, Fords, Toyotas And Rivians” 2023, <https://www.carscoops.com/2023/12/chevy-dealers-ai-chatbots-are-recommending-teslas-bmws-fords-toyotas-and-rivians/>

Test beds

Borrowing from the software industry, many companies are looking to short sprints and iteration as they work through AI's prospects and pitfalls. "There's no right answer, there's going to always be multiple answers ... start small, iterate, learn quickly, fail, get feedback, and keep moving forward," advises Melissa Pint, chief digital information officer at Frontier, a telecoms company. This careful strategy allows firms to explore potential without overextending. "Fewer than 20% of AI pilots end up in production," says Gereurd Roberts, group managing director at Seven Digital, part of Australian media company Seven West Media. "It requires a heap of experimentation in working out what's going to create value."

As Kushal Chakrabarti, chief data officer at Opendoor, an online real-estate company, notes, "if you can get a low precision prototype in a short amount of time, you can use it to answer some questions and think about the evolution, and then you can add precision and

reliability over time. It's a stronger version of working backwards." The key, say experts, is to avoid rollout for its own sake. "We see mistakes happening due to poorly defined problems and a lack of understanding of the technology without the necessary subject matter expertise," says Mr Garciga. "My fundamental advice is: be deliberate, start small, keep it scoped, and make sure you have a clear problem you're trying to solve."

Companies consider a vast array of ideas for AI applications, then carefully narrow them down to select the best use cases. "We have an intake process, by which I mean a centralised function that [considers GenAI implementation ideas from across the company, and] the number of ideas keeps increasing every day," says Mr Ulrich at Mastercard. "Through this we see ideas increasing in two ways. First, as we find what drives value, the percentage of ideas that make it through the funnel increases. And second, even if the percentage of ideas that are actually commercialised remains small, you will still see a lot happening because we're growing the denominator of ideas."

"Fewer than 20% of AI pilots end up in production. It requires a heap of experimentation in working out what's going to create value."

Gereurd Roberts, group managing director,
Seven Digital, Seven West Media



Senthil Ramani, the global lead for data and AI at Accenture, a global consultancy, says the firm sees a distinction emerging between use cases. Table stakes refer to more straightforward use cases, which include optimising customer contact centre operations, marketing functions or coding workflows. Strategic bets represent more innovative and experimental use cases.



Sandbox environments are one mechanism for AI experimentation; a space not only to test the technical aspects of an ML model for hallucination, bias and performance, but also its business impact and validity. Scott Hallworth, the chief data and analytics officer at HP, a technology company, is a strong proponent. He embraced the sandbox approach early on in the GenAI boom, where all HP employees had access to an environment where they could explore and test solutions that can make life easier for HP staff and customers. HP today supports 75 private sandboxes for a variety of uses all in the name of testing and exploring the possibilities that AI can bring to the organisation in a safe, secure and robust environment.

Key takeaways

Take a stake-based approach to experimentation. Cognisant of the risks and unknowns, firms are focusing on internal use-cases in lower-stakes domains as they ascend the maturation curve, and should focus on small, tightly scoped, problem-oriented projects.

Sandboxes and centralised oversight for pilots can help coordinate experiments. ‘Letting a hundred flowers bloom’ could impede long-term coherence and governance of AI projects. Sandboxes, testing environments and institutions to evaluate and greenlight experiments can help bring order.

In a volatile world, AI can parse signal from noise. Firms reliant on international supply chains and sensitive to economic and political headwinds are leaning on AI to inform decision-making.

Chapter two:

Rewards and returns

As AI spending soars, businesses are grappling with the challenge of balancing enthusiasm with disciplined investment strategies and realistic expectations for returns. While internal operational improvements have been the initial focus, companies are now shifting towards external AI applications that could drive significant revenue growth and unlock value from previously untapped data sources. However, measuring AI success requires patience and a nuanced approach as organisations navigate the complex landscape of long-term value creation versus short-term gains in the AI-driven future.

Rational exuberance?

Earlier in 2024 a research note from Goldman Sachs cautioned that the US\$1trn of expected spending on GenAI might risk running ahead of the revenue prospects.¹¹ The bear case included GenAI being ill suited to solving the kinds of complex problems that would justify this spend and cost. The C-suite executives interviewed for this report regard the questioning of GenAI's value proposition as welcome, and part of a natural cycle; hype followed by short-term disillusion, succeeded by proof points of value over time.

Our survey found that revenue growth has been the least important metric for AI investment to date, with only 19% of respondents agreeing that it has significantly contributed to the investment case (see figure 4).

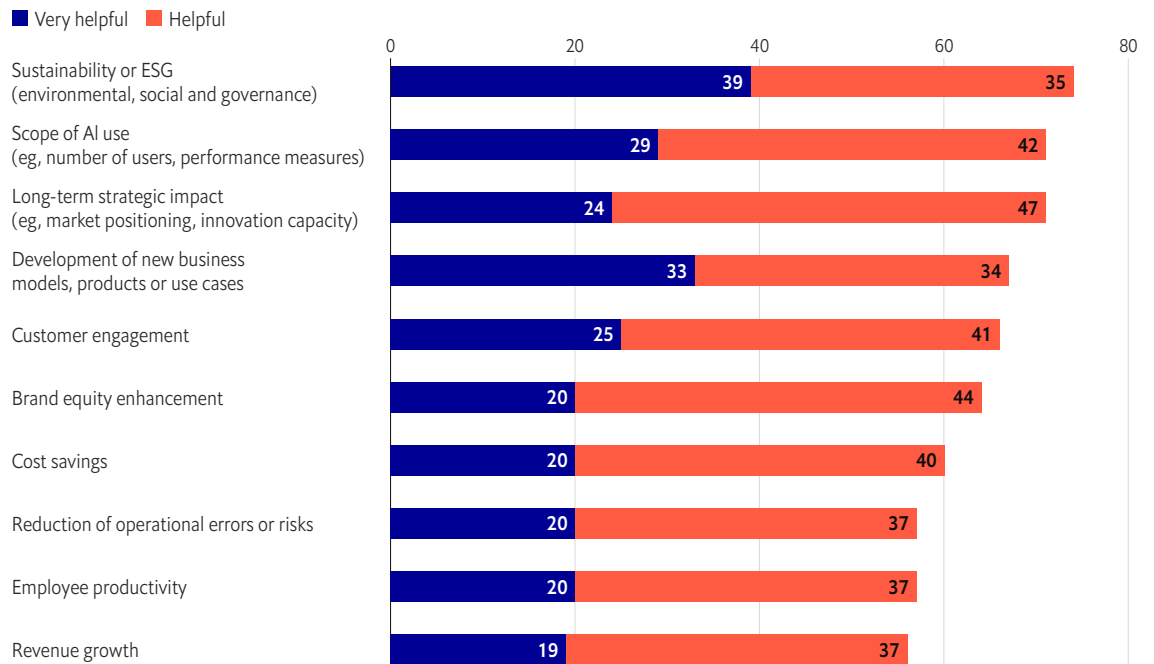


11 Goldman Sachs, "Gen AI: too much spend, too little benefit?", 2024, <https://www.goldmansachs.com/insights/top-of-mind/gen-ai-too-much-spend-too-little-benefit>

Figure 4: Revenue growth is not yet a leading driver of AI investment decisions

Executives saying the following metrics have helped their organisation build a successful business case for AI investment

Percentage of executives



Source: Economist Impact

Executives are understandably focused on the long-term view and nervous about falling prey to underinvestment. In an August analyst call, Alphabet CEO Sundar Pichai said that the risk of underinvesting in AI was “dramatically greater” than of overinvesting. Microsoft, Alphabet, Amazon and Meta Platforms are expected to incur more than US\$200bn in capital expenditure in 2024, most of which is to build AI infrastructure.¹²

But investors do expect AI to shift from a ‘tell me’ to a ‘show me’ story, with current disconnects between investments and revenue generation to come under increased scrutiny.¹³ This may be years away, not quarters, for many enterprise adopters. But vendors and firms with tangible

GenAI offerings are already posting revenue gains, as are the ‘pick and shovel’ industries like cloud computing, chips and data centres. The cloud market, for instance, could reach US\$2trn on the back of AI workflows, according to one forecast.¹⁴ Meanwhile, our survey found that six in ten enterprise architects are increasing their reliance on cloud services for AI, with a shift towards multi-cloud environments for just over a third.

Our survey also shows that deployment plans of companies will turn outwards over the next three years, shifting from internal purposes like operational improvement to external use cases (see figure 5). This may accelerate the revenue opportunity for AI as compared with using it for optimising business as usual.

12 Channel News Asia, “Nvidia fails to impress growth-hungry investors, shares fall”, 2024,

<https://www.channelnewsasia.com/business/nvidia-shares-fall-growth-hungry-investors-4573336>

13 Investors’ Business Daily, “AI Stocks: Tech Giants, Cloud Titans Face ‘Show Me’ Moment. Apple Unveils iPhone 16”, 2024,


<https://www.investors.com/news/technology/artificial-intelligence-stocks/>

14 Goldman Sachs, “Cloud revenues poised to reach \$2 trillion by 2030 amid AI rollout”, 2024,

<https://www.goldmansachs.com/insights/articles/cloud-revenues-poised-to-reach-2-trillion-by-2030-amid-ai-rollout>

Figure 5: GenAI internal and external use cases, now and in 2027

Percentage of executives




Source: Economist Impact

Indeed, many respondents (37%) are starting to see more benefits from external rather than internal use cases. This is noteworthy, given that only 24% of organisations allocate the majority of their investment to external use cases (see figure 6). It suggests that some organisations—likely those on the vanguard of expanding and scaling—are beginning to see outsized ROI from their external GenAI applications. As the remaining 88% reach this stage in the coming years, ROI may shift even more strongly towards external use cases.

Mr Ramani at Accenture highlights that productivity is a sensible rationale for initial adoption and easier to measure, but the organisations set to become tomorrow’s AI leaders will be those that transcend productivity into understanding that AI is about the “growth of revenue, new opportunities, and attracting and retaining talent”.

Figure 6: Investments versus benefits

Percentage of data scientists



Source: Economist Impact

Data intelligence

Growth and revenue will come for companies that can extract the most value from their data. Unstructured data, for instance, has historically been challenging to process, manage and use, but it offers immense value. In a medical setting, for instance, consider the ability to combine diagnostic images, text assessments, and numerical values from assessments and clinical readings. Insurance companies can use data from satellites, GPS or wearables to inform risk modelling and therefore the pricing of anything from heat waves to heart attacks.

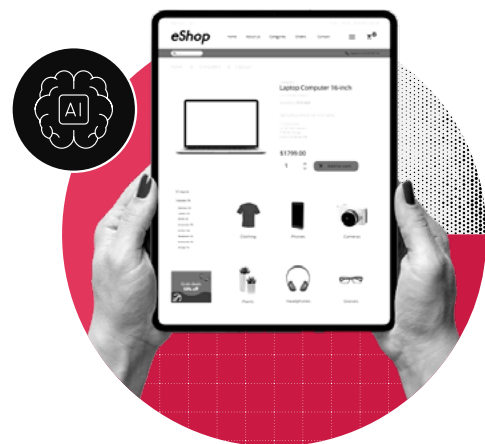
Today's AI models are appealing precisely because of how much information they can harvest and put to work. Companies report a flywheel effect: continual insights harnessed from data, using AI, in turn allows them to constantly improve products and services. For instance, Rivian is integrating its technology stack into EVs made by Volkswagen (VW), giving Rivian access to a much greater volume of data and deeper insights on usage patterns, says Mr Bensaid, the company's chief software officer. "We'll be able to bring our technology to a much broader set of vehicles and scale across different VW brands and architectures, given the inbuilt modularity and flexibility of our software architecture," he adds. This volume of data will help speed up electrification across millions of vehicles.

“The future is about really understanding your customer. The apps, products and commerce offers that they'd like to see are all becoming increasingly personalised.”

Sanjay Bhakta, chief product and technology officer, Condé Nast

This kind of data-driven approach is driving a trend of personalised products across major industries, tailored to the needs and behaviours of customers. “The future is about really understanding your customer. The apps, products and commerce offers that they'd like to see are all becoming increasingly personalised,” says Sanjay Bhakta, chief product and technology officer at media company Condé Nast.

The public sector could be a surprise winner, given that it presides over vast reserves of hitherto-neglected data. McKinsey, a consultancy, reckons that more than US\$3trn in economic value globally could be generated through the enhanced use of open data.¹⁵ When Transport for London shared data about service provision, it spurred the creation of over 600 apps in areas like journey and logistics planning, delivering economic savings of up to £95m.¹⁶ Meanwhile, the US DoD has been able to tap into its vast proprietary data to process over 10 million payment records, allowing it to identify and take action on US\$12.7bn in improper payments since 2020.¹⁷



15 McKinsey, “How government can promote open data and help unleash over \$3 trillion in economic value”, [https://www.mckinsey.com/~media/mckinsey/industries/public%20and%20social%20sector/our%20insights/how%20government%20can%20promote%20open%20data/how_govt_can_promote_open_data_and_help_unleash_over_\\$3_trillion_in_economic_value.pdf](https://www.mckinsey.com/~media/mckinsey/industries/public%20and%20social%20sector/our%20insights/how%20government%20can%20promote%20open%20data/how_govt_can_promote_open_data_and_help_unleash_over_$3_trillion_in_economic_value.pdf)

16 Deloitte, “Assessing the value of TfL's open data and digital partnerships”, 2017, <https://content.tfl.gov.uk/deloitte-report-tfl-open-data.pdf>

17 CDAO, “Advana Industry Day”, <https://storage.tradewindai.com/pdfs/Advana-Industry-Day.pdf>

“Part of the process is finding the opportunities to test these products, to get a quick read before you make a big investment. That experimentation cycle needs to be a big part of the journey.”

Jon Francis, chief data and analytics officer, General Motors

Measuring tangible gains from AI

Organisations need to develop frameworks for establishing investment returns and justifying spend. Our survey found the largest companies rely heavily on tracking key performance indicators (KPIs) and conducting post-impact or post-deployment evaluations. Smaller companies also rely on KPIs such as revenue per worker, customer retention and operational costs.

Mr Francis at GM says the company takes a highly disciplined approach, applying AI and ML where it will drive incremental revenue or mitigate costs. “It cannot be technology for the sake of technology; it has to be tied to value and impact for the organisation.” The Detroit-based carmaker collects data on small-scale pilots and proof of concept projects. “Part of the process is finding the opportunities to test these products, to get a quick read before you make a big investment. That experimentation cycle needs to be a big part of the journey,” adds Mr Francis. But the company remains laser-focused on business outcomes. “Against every dollar, every penny of compute, we should be able to articulate what the return is.”

Mahindra Group is focusing on challenges that go beyond basic productivity improvements. As Mr Kapoor at Mahindra states, the organisation is “focused on hard things to solve” particularly on the factory shop floors, to reduce mean time to repair and mean time between failures or to improve customer experience. The latter includes innovations like GenAI-powered assistance for faster car servicing. Condé Nast determines ROI across a checklist of outcomes including lowering costs, increasing speed and efficiency, improving audience experience, or increasing time on site. Unilever allocates its spending across AI for productivity, creativity and growth. Seven Digital has a similar approach. “We are taking a practical approach in using AI to improve the business model—efficiency, outcomes and upskilling, which is organisational transformation,” says the group managing director, Mr Roberts.

Siam Commercial Bank’s former chief digital banking officer, Chalee Asavathiratham, emphasised the need for AI teams to work with finance to measure the outcomes and ROI to avoid getting lured by novelty. He warned that AI teams may “neglect to ask what kind of numbers, what kind of ROI” an AI project would generate. “[AI teams] would sometimes get lured [by] the coolness of the projects,” he explains.

Frontier's Ms Pint makes a similar argument. "There is no such thing as a technology strategy. There's only a business strategy that technology supports. So whatever you start doing, understand the business strategy very, very well, and then tie your technology strategy to it."

"There is no such thing as a technology strategy. There's only a business strategy that technology supports. So whatever you start doing, understand the business strategy very, very well, and then tie your technology strategy to it."

Melissa Pint, chief digital information officer, Frontier

At Opendoor, executives determined that productivity should not be the central ROI measure, instead focusing on top-line growth and improved consumer experience. While the company's use of AI yields productivity gains, these are not at the forefront of their approach. "Our priority is better use of human intuition and improved consistency, which improves customer experience," says Mr Chakrabarti at Opendoor. "It's better for employees, customers and the business."

Building a strong data foundation can help companies pursue the higher value goals of revenue and business model innovation. "We got started with AI by building a strong data foundation, then we pivoted to a commercial use case around variable pricing of select add-on products," recalls Ms Clements at JetBlue. "We first used the model in a very narrow set of markets, and we were immediately struck by the incredible value that it delivered right out of the gate. This use case was the catalyst for us to understand the power of AI and the ways it could transform our business; it got the creative juices flowing across the organisation to think about leveraging the technology to drive even more value."

Cost control is a valid goal too. For companies with significant operational expenditure, even optimising spending by a few percent could lead to savings of millions, some of our executive interviewees have observed, and these operational savings can be re-invested in innovation.





“It’s very hard to institutionalise AI technology and data analysis with employees.

You cannot just measure over the short-term; you need to think over a three-year timeline.

Helen Choi, chief digital and information officer, CJ CheilJedang

The virtue of patience

Defining ROI requires identifying objectives and establishing realistic timelines for achieving them, calibrated based on an accurate understanding of AI capabilities, implementation challenges, and organisational readiness. When it comes to digital transformation and AI, many executives expect short-term performances and results, explains Helen Choi, chief digital and information officer at CJ CheilJedang, a South Korean consumer goods firm. “But it’s very hard to institutionalise AI technology and data analysis with employees. You cannot just measure over the short-term; you need to think over a three-year timeline. This is because it is not just a matter of data or AI technology implementation,” says Ms Choi. “Achieving real business value, led by digital technology, involves a long journey of innovation across ways of working and organisational change management.”

“A significant challenge lies in the fact that the monetisation and business models for value creation through GenAI are still not well established,” says Takaaki Sato, the senior executive vice president and chief technology officer at NTT Docomo, a Japanese telecoms company. “There is uncertainty about whether to invest in general-purpose LLMs or specialised smaller LLMs, [and] unclear ROI at this stage.”

While a laser-focus on returns can mitigate excessive spending, companies might take one step back to take two steps forward. For example, some of the algorithms currently used by Dream Sports, an Indian fantasy sports company, produced negative business metrics initially. “Obviously, the product teams and the finance teams got a little nervous about that. [It’s about] being able to persevere and persist and educate them,” recalls Amit Sharma, the company’s chief technology officer. In regulated sectors like pharmaceuticals and biotech, meanwhile, achieving quick wins can

be challenging, as it takes years to bring a pharmaceutical innovation from the lab to the patient. “AI accelerates the process of identifying synthesisable molecules and advancing them to phase two or three trials,” notes Mr Bucher at Novartis. “However, it may take five to ten years to determine AI’s effectiveness, once the product’s safety is confirmed.”

The rapid rate of AI progress and innovation might also merit strategic patience, with some experts likening it to waiting for Google to emerge from the many search engines in the 1990s (it appeared six years after Archie, the debut product).¹⁸

Jeff Martin, senior vice president and chief data officer at TD, cautions against “being too confident [about] where the future is going”. He adds that companies should assume they will need to pivot on every decision they’re making, “so be sure to be ready, flexible and agile”.

“AI accelerates the process of identifying synthesisable molecules and advancing them to phase two or three trials. However, it may take five to ten years to determine AI’s effectiveness, once the product’s safety is confirmed.”

Bernd Bucher, chief information officer, Novartis

Key takeaways

Productivity and efficiency gains are early, measurable wins, but true leaders will seek more ambitious returns. Revenue, talent attraction, and novel products and services are among the loftier targets. Some firms can also unlock value by using AI to extract and apply their proprietary data in new ways, including developing more personalised products.

Organisations must strike a fine balance between business returns and the time to achieve true value. Companies will soon need to demonstrate financial returns from their AI investments, but their timelines must strike the right balance. Experimentation takes time, and some projects will inevitably fail to deliver. Firms need an initial runway to upgrade their data infrastructure. Regulatory constraints also limit the speed of deployment in sectors like healthcare. There is also merit in watchful waiting as the AI vendor and product landscape evolves.

18 SEO Mechanic, “The Complete History of Search Engines”, 2023, <https://www.seomechanic.com/complete-history-search-engines/>

Chapter three:

The infrastructure refit

As organisations harness AI's potential, they face a critical challenge: their data infrastructure is woefully unprepared, akin to having Victorian-era plumbing in a modern building. Silos, latency and security are all inhibitors for AI deployment. From accessing valuable but unstructured data to navigating the complexities of cloud solutions, organisations need to reimagine their data architecture.

One of the biggest shifts in the AI era is not the power of the models themselves, but the underlying plumbing needed to deliver them at scale. Many of the world's largest companies are sitting on the technological equivalent of Victorian-era pipes that cannot handle the demands of AI, with legacy systems that are creaking at the seams.

Only 22% of organisations in our survey say their current architecture is fully capable of supporting the unique demands of AI workloads, and just 23% say their current architecture fully integrates AI applications to relevant business data. Even among the largest companies we surveyed, these rates only rise to 28% and 27%, respectively. These deficits hinder progress and likely contribute to quality concerns, with only 37% of executives and even fewer practitioners (29%) believing that GenAI applications at their organisations are production ready.

The demands of AI and advanced analytics workloads were already creating infrastructure headaches before GenAI, largely because of the different types of data generated by companies and the variety of storage mechanisms developed to handle them.

Many companies used cloud-based warehouses to store highly structured data that have clear classifications and labels, managed by expert IT teams and easily queried. However, as much as 90% of data collected by organisations is unstructured and, therefore, unfit for warehousing.¹⁹ As ML and AI advanced, unstructured data came to be seen as ever more valuable because of ML systems' ability to spot patterns in vast troves of this previously neglected data.

Only 22% of organisations in our survey say their current architecture is fully capable of supporting the unique demands of AI workloads.

¹⁹ MIT Sloan, "Tapping the power of unstructured data", 2021, <https://mitsloan.mit.edu/ideas-made-to-matter/tapping-power-unstructured-data>

Data lakes, capable of storing vast quantities of both structured and unstructured data, emerged as an alternative, allowing organisations to leverage unstructured data, but these led to fragmentation between lakes and warehouses. In effect, this led to silos, patchy support for use cases, and incompatible security models, all of which made it harder for companies to utilise their data effectively. Then, over time, a new model called the data lakehouse emerged, aimed at addressing these problems. This model combines aspects of both lakes and warehouses and is built on open-source software, using open standards.

Looking ahead, speed is the primary architecture constraint that organisations plan to address, with real-time data processing voted the top capacity gap by 47% of respondents, followed by robust data pipelines and security (see figure 7). Architects' views were also corroborated by data engineers, who were the most likely to rank streaming data issues as their top challenge in maintaining data quality.


Real-time streaming of live production data is crucial for AI applications that require immediate querying. Mahindra Group, for instance, incorporates real-time data on solar panel performance into an algorithm that directs maintenance and cleaning crews around their renewable energy installations for effective cleaning and maintenance schedules. Similarly Dream Sports' use of AI to decide when to generate lucrative contests around sports events depends on knowing exactly how many users are building a fantasy team for a specific event at any moment.

Diseconomies of scale

Finding the right infrastructure can become more difficult as companies get bigger. Diversified companies with individual business units, and conglomerates comprising entirely independent companies, must determine where the data go. For instance, having data centralised in company-wide stores allows insights to be drawn across the group, while opting for individual units to manage the data gives more flexibility to analytics and data science teams to work with it in ways that suit their own needs.

Figure 7: Architecture limitations

The most significant gaps in respondents' current architecture that need to be addressed in the next 3-5 years
Percentage of enterprise architects



Source: Economist Impact

At Mahindra Group, core functions like cloud and cybersecurity are managed in common, but individual architecture decisions are left to business units. The firm takes a ‘problem first, solution later’ approach, emphasising collaboration between business and technology teams. Its car company uses a data mesh to store driver relevant information such as journey history separately from manufacturing data, while allowing the two datasets to be combined if needed. This reduces latency for driver queries, because only relevant information is searched, but allows for supply chain problems to be identified—for example, if several drivers report a common problem the company can trace the manufacturing history of the relevant component in each car. Mahindra’s finance team, meanwhile, uses a more flexible architecture so that it can consider personal loans, fixed deposits, insurance, mortgages and loans for small and medium-sized enterprises simultaneously in order to assess risk.



It would be wrong to assume that start-ups or digital native companies are free from infrastructure challenges. Dream Sports, for example, was still querying production databases for its reporting needs several years after its founding and only set up a data store once the company had found product-market fit. However, such companies show a remarkable willingness to reengineer their data architecture from scratch when needed, perhaps because they lack the substantial technical debt.

Rivian, an EV manufacturer, has integrated information from across its complex systems, including supply chain, commercial operations and financial departments, into a unified environment, enabling data-driven decision-making at all levels of the organisation. This benefits not just AI and GenAI applications, but also the entire gamut of data-linked business processes from dashboards and analytics to customer profiles and history.

GenAI can pose an infrastructure pressure for some because of the demands of new use cases, with their corresponding unique tooling requirements. One of these is Retrieval Augmented Generation, or RAG, the process through which an LLM checks the answer it has generated against an external source, like an encyclopaedia, and then refines it. When building their own LLMs, organisations supply internal data, such as customer support interactions, as the authoritative external source of information. But before this can happen the internal data must be cleaned and prepared for ML, which requires specific tools and technologies.

In their rush to embrace GenAI, many companies are building and maintaining duplicates of data or systems to clean and prepare data, for example, by adding an additional data pipeline to the company’s infrastructure. This adds further complexity to already complicated data infrastructure.



“You're kind of building in a vacuum. And every application is building its own set of infrastructure and tooling around large language model operations or machine learning operations,” says Ms Vaezy at Providence.

Indeed, in a world of GenAI, data infrastructure could become congested, overlapping and harder to keep track of. When asked what would yield the biggest benefits for their productivity, data engineers identified ‘simplifying data source connections for ingesting data’, ‘using a single unified solution instead of multiple tools’ and ‘better visibility into data pipelines to find and fix issues’ among the top interventions.

Implementing traditional AI use-cases prematurely, without a solid data foundation, could expose companies to operational, security and reputational risks. “Back in March of 2023, when all the excitement [around GenAI] was taking place, I knew that we were going to run into a problem that everybody was going to be excitedly jumping into things and potentially exposing the company in ways we shouldn't be exposed,” says Mr Hallworth at HP.

Getting data right—and getting the right data

Many organisations are increasing their reliance on cloud services to support AI workloads, as reported by 61% of enterprise architects in our survey, with a trend towards multi-cloud solutions among 37%, based on factors including control, security, data ownership, compliance and cost.

Companies in our survey often report moving some workloads from public to private cloud environments since first investing in AI, where resources like server hardware are dedicated to a specific company rather than serving many clients at once. The control, security, data ownership and compliance aspects of private cloud infrastructure can make it an attractive option for integrating GenAI applications. Compliance needs and the sensitive data required to train and run models often demand tailored infrastructure.

On the other hand, the scalability and flexibility offered by public cloud systems at low upfront costs are valuable. Shorter development cycles, integrated technical frameworks, experimentation at scale and the ability to rapidly scale can make public cloud systems a tool of choice. But there is a balance between cost-efficiency and control, security and flexibility that remains an open question.

“You can have all the AI in the world, but if it's on a shaky data foundation, then it's not going to bring you any value.”

Carol Clements, chief digital and technology officer, JetBlue


Beyond the question moving data into storage comes the knotty issue of its accuracy and security, along with the appropriate access controls. Incorrect or conflicting data, after all, could lead a GenAI model to offer incorrect information, and weak controls could lead to compliance breaches. As Ms Clements at JetBlue said: “You can have all the AI in the world, but if it's on a shaky data foundation, then it's not going to bring you any value.”

“Ultimately, all of this runs on data. If your data isn't properly categorised, it's going to be very difficult to benefit from any of this technology, so you have to focus on the fundamentals of good data engineering and data management,” argues Mr Vohs at Molson Coors. When it comes to these fundamentals, enterprise architects identified security, data privacy and data silos as the three biggest challenges faced when managing data for AI applications (see figure 8).

Figure 8: The challenges of data gatekeeping

Enterprise architects' biggest challenges in managing and controlling data for AI applications

Percentage of enterprise architects




Source: Economist Impact

One challenge introduced by GenAI, in particular, is how much information to put into the cloud for analysis. Some are reluctant to push more data into the cloud without a clear use case in mind. This reluctance is pronounced in sectors like energy and banking, which have already accrued substantial amounts of technical debt and data silos through legacy systems and mergers and acquisitions, for example. Both Spanish energy company Repsol and Canadian bank TD also emphasise the potential expense of excessive data upload, both directly in cloud-related costs, and indirectly in terms of requiring staff and teams to prepare data and build and maintain pipelines to push data into storage.

While companies in the past tended to ignore unstructured data—despite it making up the vast majority of what companies hold—it has taken on increasing importance as AI tools make it easier to access.²⁰ UPS uses a data mesh to classify data and analytics tools that allow both data scientists and business teams to engage in machine learning operations (MLOps). “It is powerful because it can combine different ML models in a quick time frame. If tomorrow we wanted to infuse a GenAI model into it, we would be able to do that because we’ve got all the platforms and orchestration in place to do it,” says Mr Finnerty at UPS.

Figure 9: GenAI quality: leaders and laggards
Strategies employed for evaluating the quality of GenAI responses
 Percentage of data scientists



Note: for laggards, n=50; for leaders, n=37; insights are only directional and suggestive
 Source: Economist Impact

20 MIT Sloan, “Tapping the power of unstructured data”, 2021, <https://mitsloan.mit.edu/ideas-made-to-matter/tapping-power-unstructured-data>

The burden of housekeeping

GenAI stands out among other uses of data in the technology arsenal for how much upkeep and maintenance it requires. LLMs, for example, suffer from so-called model drift over time, as the information they were trained on becomes dated and their performance deteriorates.²¹ Indeed, Mr Hallworth from HP likens AI models to cars that start to depreciate the moment they are driven off a parking lot.

Upkeep comes in at least three forms. First, datasets must be constantly validated and checked for completeness, recency and validity. These checks can be performed by open-source tools and application programming interfaces (APIs) integrated into the existing data sets and clouds. Quality assurance can also be applied across data pipelines to prevent faulty data from

entering the organisation's systems in the first place. Very large firms (revenue over US\$10bn) are more likely than others to use ML to detect data anomalies, our survey found. And when it comes to data upkeep for GenAI models specifically, 'leaders'—companies with production ready GenAI tools—are also more likely than others to use AI to assess the data (see figure 9).

Second, for GenAI systems, the encyclopaedia-like external data sources used by LLMs to refine their answers must be regularly updated to maintain the quality of information they contain.

Finally, the models themselves must be constantly tested for efficacy. One example highlighted in figure 9 is automated evaluation collection, whereby ML algorithms are used to quantify the accuracy of outputs generated by an LLM in comparison to expected answers.²²

Key takeaways

Only a fifth of organisations have an architecture fully capable of supporting AI workloads.

This is a major constraint to operationalising AI. Real-time data processing is one high-value capability sought by companies, which can only be delivered with an infrastructure refit. Unified infrastructure helps AI deliver on its potential not just as a productivity tool, but also as an input for better decision-making at the speed businesses need.

Size or digital savvy is no predictor of success in infrastructure transformation.

Large companies have more complex decisions to make around centralisation and devolution, and digital native firms may have more complex tooling environments due to early adoption of best-of-breed software.

AI needs housekeeping. Datasets must be constantly validated and checked for completeness, recency and validity, performed by open-source tools and APIs integrated into the existing data sets and clouds. Quality assurance should be applied across data pipelines to prevent faulty data from entering the organisation's systems in the first place. External data sources used by LLMs to refine their answers must be regularly updated to maintain the quality of information they contain. And the models themselves must be constantly tested for efficacy.

²¹ IBM, "What Is Model Drift?", 2024, <https://www.ibm.com/topics/model-drift>

²² Microsoft, "Evaluation of generative AI applications", 2024, <https://learn.microsoft.com/en-us/azure/ai-studio/concepts/evaluation-approach-gen-ai>

Practitioner perspectives

Enterprise architects: walking the tightrope

The role of a data and AI architect is complex, challenging and crucial in today's technology-driven organisations. Gregor Hohpe, who has held leadership positions at tech giants like Google and Amazon Web Services, defined a software architect as "a technical leader who thrives on complexity but insists on simplicity". Architects are charged with the design, implementation, and oversight of a scalable and secure data and AI architecture, requiring a deep understanding of current and emerging technologies.

Our survey found that just 22% of organisations have architectures that fully support AI workloads without modifications, and only 23% can connect AI applications to relevant business data without changes. Architects are the ones now tasked with bridging this gap and helping organisations achieve comprehensive AI integration. Key capability gaps include real-time data processing, identified by nearly half (47%)

of architects as a significant need. Other concerns include data privacy and security breaches. To strengthen governance, architects are focusing on implementing custom APIs (58%), human-in-the-loop processes (55%) and integrating ML models into familiar data tools (52%). This approach reflects the need for both technical solutions and human oversight in AI implementation, ensuring that AI capabilities are seamlessly woven into the fabric of organisational operations.

Technologies that enhance accessibility and efficiency across different platforms are becoming increasingly important. Natural language processing (NLP) is at the forefront of this trend, with 58% of architects believing it will become the primary (or only) tool in data pipeline creation. Sixty-seven percent say the same about its impact on workflow development, as do 52% about its role in report generation. This indicates a shift towards more intuitive, user-friendly tools in data and AI operations.

In the current economic climate, with mounting pressure to reduce the total cost of ownership, these cross-platform capabilities are becoming more critical. The emphasis on accessibility and efficiency not only reduces costs but also promotes the democratisation of data and AI, enabling a broader range of users to access and leverage these powerful tools and insights across the organisation.

The challenges faced by data and AI architects are numerous, but so are the opportunities. As they seek to develop strategies for implementing ethical AI, managing data ecosystems and democratising data access, architects will face a delicate balancing act of innovation, security and accessibility.

Data engineers: bringing order to the data storm

Data engineering is under immense pressure to meet the evolving demands of modern enterprises, with two dominant challenges—fostering high-data quality for end users and ensuring security and governance. These objectives come on top of data engineers' responsibility for managing most aspects of the data engineering lifecycle.

The need for unification, simplification and democratisation across the data engineering discipline is at a fever pitch. Today, nearly two-thirds of organisations are fully dependent on data engineers for every aspect of data pipeline creation and management, and almost half of data engineers spend most of their time configuring and fixing data source connections. Data engineers face significant burnout and fatigue given the number of data sources and tools required to configure usable data pipelines.

The emergence of NLP as a tool for democratising data pipeline and workflow creation is a promising trend, with more than 97% of enterprise architects predicting that NLP will either supplement or replace traditional methods for pipeline creation and workflow orchestration. This shift means, in the future, that data engineers will increasingly enjoy support from intelligent automation to handle the complexity of real-time data processing and pipeline management, allowing them to focus on more strategic tasks.

At the top of the wish list for most data and AI teams are simplified data source connections, integrated GenAI tools for coding assistance, and unified platforms to reduce the reliance on disparate systems. Architects are placing significant bets on real-time data processing and robust data pipelines, with nearly 75% identifying one of these as a critical goal for their organisations' data and AI capabilities. As streaming architectures become more prevalent, with 46% of companies migrating a significant number of batch pipelines to real-time streams in the last year alone, the emphasis on latency and real-time insights is intensifying.

In sum, data engineering is at a critical juncture. The field is moving towards greater efficiency through automation and real-time data processing, but it remains anchored by the foundational challenges of security, governance and data quality.

Chapter four:

Find your edge

As AI becomes a baseline for business, companies will need to find their unique path by integrating foundation models with proprietary data and domain expertise, balancing off-the-shelf convenience with the power of tailored solutions, all while safeguarding their most valuable asset: their data.

As foundation models become commoditised, usage will be a baseline, not a boost. An oft-repeated warning in the contemporary discourse is that AI will not replace humans, but humans using AI will replace those that don't. The same goes for companies. Each must find their edge—their unique recipe of data, need and know-how. “You need to build on top of your existing value. In our case it is our medical knowledge, user experience, connection to wearables, and using GenAI to produce helpful insights or propose answers to questions,” says Mr Bugaev at Flo Health.

Mix and match

Our survey found that two-thirds of organisations see significant potential in integrating GenAI models with their own proprietary data. Most will mix and match based on workflows and needs. For instance, 58% of data scientists augment their LLMs with proprietary data through a process known as RAG—the use of a knowledgeable external source by an LLM when inferring or producing an output—while 45% say they use LLMs off the shelf without connecting to their data (see figure 10). Meanwhile, 21% do both.


Mr Hallworth at HP draws a pizza analogy to describe the difference. Off-the-shelf AI models may be convenient and time-saving, but a home-cooked one that takes longer to make tastes better. A ‘supermarket pizza’ AI equivalent might write first draft marketing copy or meeting summaries, while advanced tasks like customer-facing interactions, product insights or forecasting need to be cooked at home.



Figure 10: Building and enhancing LLMs with enterprise data

Nature of current GenAI projects among data scientists

Percentage of data scientists



Source: Economist Impact

“Most of the value that you see on AI is buried in the bottom,” explains Accenture’s Mr Ramani. “The top is immediate benefits like productivity, but that’s the starting point. You’ve not gone into your core value chain, that’s when you go beneath the iceberg. That is where your proprietary data is.”

“Most of the value that you see on AI is buried in the bottom. The top is immediate benefits like productivity, but that’s the starting point. You’ve not gone into your core value chain, that’s when you go beneath the iceberg. That is where your proprietary data is.”

Senthil Ramani, global lead, data and AI, Accenture

Foundation models are the coffee grinder; the beans are the company’s data. Using that data as part of a ‘compound’²³ system helps organisations brew AI results that consider their unique domain or situation. Blending multiple models to meet varied use cases is a common trend, our survey shows (see figure 12), as is integrating proprietary data into GenAI models, which most respondents are either doing or exploring.


Respondents from the financial services industry were by far the most likely to recognise significant potential in integrating GenAI models with proprietary data, with 80% agreeing (see figure 11). Across sectors, perspectives on this differ substantially, reflecting diverse views on the value of organisational data, the complexity of processing it, the risks of using it in GenAI models and the cultural or operational readiness to implement such technologies.

23 Bair, “The Shift from Models to Compound AI Systems”, 2024, <https://bair.berkeley.edu/blog/2024/02/18/compound-ai-systems/>

Figure 11: The upside of proprietary data across industries

Respondents who agree/disagree that their organisation sees significant potential in integrating GenAI models with its own proprietary data

Percentage of all respondents



Source: Economist Impact

Choosing models

Model selection is an important means of tailoring AI to each organisation’s unique assets. Open-source AI models, for instance, allow anyone to ‘study, modify, use and share’ them.²⁴ Although there are debates over the definition of open source for AI specifically, such as whether it should include the data used to create models, the presence of open-source tools has widened access to AI.²⁵ Meta’s Llama, for instance, has been downloaded 300 million times, allowing organisations to build AI models suited to their purposes.^{26,27}

A hybrid approach of utilising both open- and closed-source AI is a common strategy, our survey found, with 75% of organisations employing open- and closed-source models based on industry-specific factors, and 89% will do so by 2027. Closed-source models are typically deployed for internal applications, while open-source models are increasingly used by those expanding into external use-cases.

Organisations cite several reasons for open-source adoption. Greater development transparency tops the list at 52%, followed by access to open-source community resources, the opportunity for proprietary intellectual property (IP) development, mitigation of vendor lock-in, and greater customisation, privacy and security control. This underscores the flexibility and control benefits offered by open-source models.

24 Open Source Initiative, “The Open Source AI Definition – draft v.0.0.8”, <https://opensource.org/deepdive/drafts/the-open-source-ai-definition-draft-v-0-0-8>

25 Economist Impact, “Open sourcing the AI revolution”, 2024, <https://impact.economist.com/perspectives/technology-innovation/open-sourcing-ai-revolution>

26 The Economist, “Meta is accused of ‘bullying’ the open-source community”, 2024, <https://www.economist.com/business/2024/08/28/meta-is-accused-of-bullying-the-open-source-community>

27 Microsoft, “Tiny but mighty: The Phi-3 small language models with big potential”, 2024, <https://news.microsoft.com/source/features/ai/the-phi-3-small-language-models-with-big-potential/>

But closed-source models offer some advantages. Superior performance was a leading reason to choose closed-source models at 59%, followed by user-friendly interfaces and documentation, exclusive features, and seamless integration with existing proprietary systems. Some also find that open-source models require more in-house expertise and skills to adapt and develop compared with off-the-shelf closed-source models.²⁸

Industry leaders are developing frameworks to inform their model decisions. Accenture helps organisations through what it calls its switchboard, a proprietary tool set and framework, enabling organisations to decide which models to use based on the business use case, relevance, cost, accuracy and latency. These dimensions help determine whether an organisation needs open source, frontier or other models.

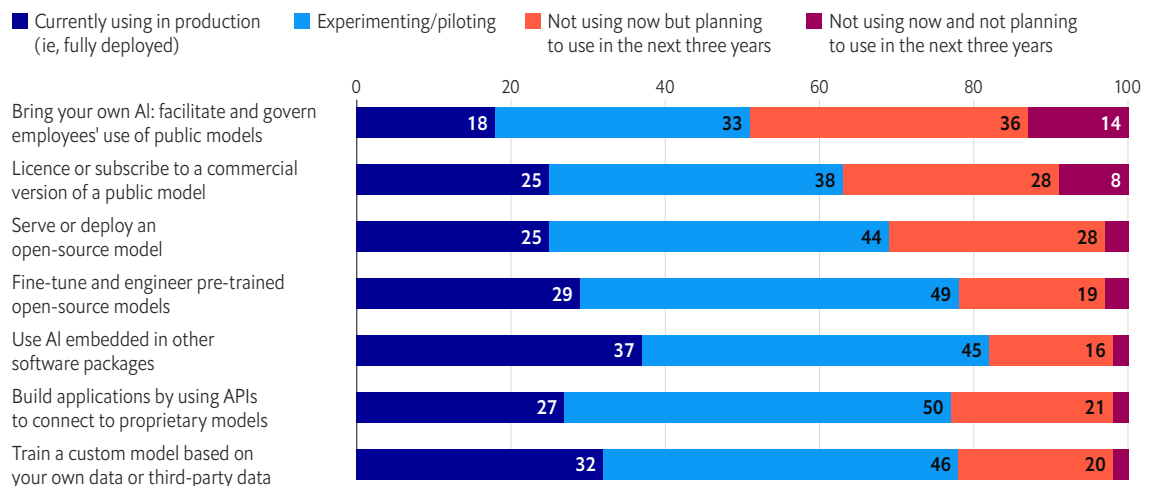
Relatively few companies see merit in developing their own LLM from scratch.

“We don’t believe a company of our size can promptly build something truly meaningful in terms of home-grown LLMs,” says Mr Bugaev at Flo Health. Equally, companies will not be able to find a measurable advantage relying solely on off-the-shelf models.

“We believe that partnerships between tech companies and pharma/healthcare are essential,” says Mr Bucher at Novartis. “Tech companies bring strong platforms, data capabilities and AI talent, which are invaluable in creating usable models. Pharma companies possess deep expertise in safety, clinical trials and production, which are critical for developing potential drugs. AI serves as an accelerator, enhancing our combined efforts. By leveraging our distinct strengths, we can achieve greater advancements together.”

Figure 12: Organisations are blending a variety of GenAI models

Adoption of GenAI models, by type
Percentage of executives



Source: Economist Impact

28 Economist Impact, “Open sourcing the AI revolution”, 2024, <https://impact.economist.com/perspectives/technology-innovation/open-sourcing-ai-revolution>


Companies and organisations have access to the best fuel for AI models: the right and most relevant data, not merely data sourced from the open internet. Dream Sports, for instance, found third-party LLMs ineffective, even if they were using the latest application, because the company's language is too specific. Models have to be trained on the company's data.

Higher quality internal data improves the efficiency and performance of models. For instance, Rakuten built a deep learning foundation and embeddings using far less processing power because it trains them based on transaction data from its e-commerce site, according to the company's chief AI and data officer, Ting Cai. "We can encode more information in fewer dimensions because we have high-quality transaction data that better indicates our users' interests, and using this data allows us to create more efficient models."

"We can encode more information in fewer dimensions because we have high-quality transaction data that better indicates our users' interests, and using this data allows us to create more efficient models."

Ting Cai, chief AI and data officer, Rakuten Group

Shell is also adapting foundational models to work on internal data, for example, customer interactions or past contracts. "I'm not too worried about hallucination in the language models Shell uses because of the controls we have in place before deployment. We restrict these models to work only with our controlled data, which has gone through the necessary rigour, including the right human supervision, before becoming a part of our work processes," says Robbert Van Rutten, chief information officer of Shell, the energy company. "We are also working on developing our own models by pre-training or fine-tuning LLMs to further minimise this risk."



“In the future, the market will move to the use of small private LLMs. But GenAI is a new technology. We still need to learn how to launch and put initiatives into production in an agile way.”

Juan Jose Casado, chief digital officer, Repsol

Companies can, of course, use third-party models for initial tinkering before deciding where to invest more technical resources. At Repsol, for example, the company fed internal data into existing LLMs offered by companies like OpenAI, Microsoft and Google rather than building a GenAI pipeline from scratch as it experiments with use cases and seeks to explore the technology’s potential. “In the future, the market will move to the use of small private LLMs,” predicts Juan Jose Casado, chief digital officer at the Spanish energy company. “But GenAI is a new technology. We still need to learn how to launch and put initiatives into production in an agile way.”

Home-grown models are also appealing because they put companies in control of the data. Many are worried about data leakage in closed-source LLMs, such as sensitive or restricted material being dropped into external models. “We have strong views on owning and protecting our data, so we’re very cautious about where we let our data go,” says Ms Clements at JetBlue. “We don’t let it leave JetBlue.”

Key takeaways

To unlock value and find their edge, companies will mix and match off-the-shelf AI with open-source and in-house models. AI will soon be a baseline capability; the organisations that excel will mix and match models to complement their unique data and know-how. Only they have access to the best fuel for AI models: the right and most relevant data.

AI promises to unlock proprietary data that was hitherto hidden or dormant. This will drive higher-performing models through accuracy, efficiency and relevance. Off-the-shelf AI will be useful for low-cost tinkering and lower-value outputs. For true value creation, and for greater control and security, companies will combine open- and closed-source models, while larger firms may see value in developing their own models from scratch.

Chapter five:

Guardrails and governance

As AI becomes ubiquitous, organisations face a complex web of risks, from data breaches to regulatory compliance. Companies are developing multifaceted governance approaches, from institutional measures such as AI centres of excellence, principle-based approaches like human-in-the-loop guardrails, and deployment of innovative technical solutions like synthetic data and RAG to improve AI reliability. The challenge lies in striking the delicate balance between harnessing AI's transformative power and maintaining robust safeguards, all while adapting to an ever-evolving regulatory environment and calibrating human judgement and machine intelligence.

Data security and governance had already become board-level priorities over the last decade as cybersecurity attacks—and the regulatory punishments—have grown more severe. A flurry of regulations like the Digital Operational Resilience Act in Europe are also raising the stakes of governance failures. Our survey underscores these priorities, finding that data privacy and security breaches are the top concern for 53% of enterprise architects, while security and governance are the most challenging aspects of data engineering for engineers (cited by 50% as a top-three challenge).

Data privacy and security breaches are the top concern for 53% of enterprise architects, while security and governance are the most challenging aspects of data engineering for engineers.

“Numerous issues, such as privacy protection, security measures, IP issues, ethical concerns and the problem of deepfakes need to be addressed,” says Mr Sato of NTT Docomo. “It is essential to maximise the benefits of AI technology by overcoming these challenges in collaboration with relevant parties,” he adds.

Worldwide, legislation is ratcheting up. One 2024 tracker estimated a total of 762 AI legislation documents across 45 states in the US alone.²⁹ This poses a challenge for companies that need to experiment with new systems while the rules governing their use are in constant flux.

According to Ian Botts, chief technology officer of Fanatics Betting & Gaming, effective AI governance at an organisation encompasses three components. “First, it's building the guardrails for what you can and can't do, so that personally identifiable information isn't used inside of a model to specifically target or take advantage of a particular human. Second, it's mechanisms—and


²⁹ Regulatory Transparency Project, “A Sensible Approach to State AI Policy”, https://rtp.fedsoc.org/blog/a-sensible-approach-to-state-ai-policy/?utm_content=311457390&utm_medium=social&utm_source=twitter&hss_channel=tw-574405888

building out mechanisms—that enforce those guardrails. And third, it's having dedicated data governance, which is almost like a compliance function to ensure that it's being audited and that [companies are] being transparent. I think the transparency point is incredibly important because it drives better decisions.”

Data innovations can help companies tackle governance risks like bias or unreliability. For example, some organisations may lack access to the necessary data to build what they need, in turn limiting the quality of their AI systems. Synthetic data, generated using a purpose-built mathematical model or algorithm, can be brought in to help solve these data science tasks.³⁰ It can improve performance and reliability by providing a larger volume of more diverse data for training, such as including underrepresented data classes or categories, for instance. Closing gender gaps and eliminating racial biases are promising areas where synthetic data can make a difference.³¹ It can also lower data acquisition costs, eliminate data bottlenecks and preserve data privacy. While it has risks and downsides, with the right guardrails, synthetic data could help improve the quality of AI outputs, but it remains an untapped resource, according to our survey, with only 7.5% of data scientists utilising it in their GenAI projects.

RAG is another useful technique to control quality and enrich knowledge. “With this technique, Rakuten’s GenAI models retrieve real-time inventory and pricing information that they can use to incorporate the latest information or enterprise knowledge into applications such as search, recommendation systems, and adverts,” says Mr Cai, the company’s chief AI and data officer.

Ian Botts at Fanatics also explains how its AI and GenAI models are being used to detect and enhance governance issues related to responsible gaming. “We want your bet to be additive to your fandom,” he says. “If [our models] identify someone and say with very high accuracy that there’s a problem that spend is exceeding their means, or they’re showing troubling behaviours or patterns, [we] could start to nudge them into free-to-play games, lower cost things, try to reduce that burn or encourage them to take a timeout. Present them with information. It’s worth the additional investment ... to get someone back into a position that’s healthy and sustainable for them. AI can be used to spot and understand problems, and GenAI can be used to understand and influence [in a personalised way].”



- 30 The Alan Turing Institute and The Royal Society, “Synthetic data - what, why and how?”, https://royalsociety.org/-/media/policy/projects/privacy-enhancing-technologies/Synthetic_Data_Survey-24.pdf
- 31 UNU Macau, “Bridging the Gender Data Gap: Harnessing Synthetic Data for Inclusive AI”, 2024, <https://unu.edu/macau/blog-post/bridging-gender-data-gap-harnessing-synthetic-data-inclusive-ai>

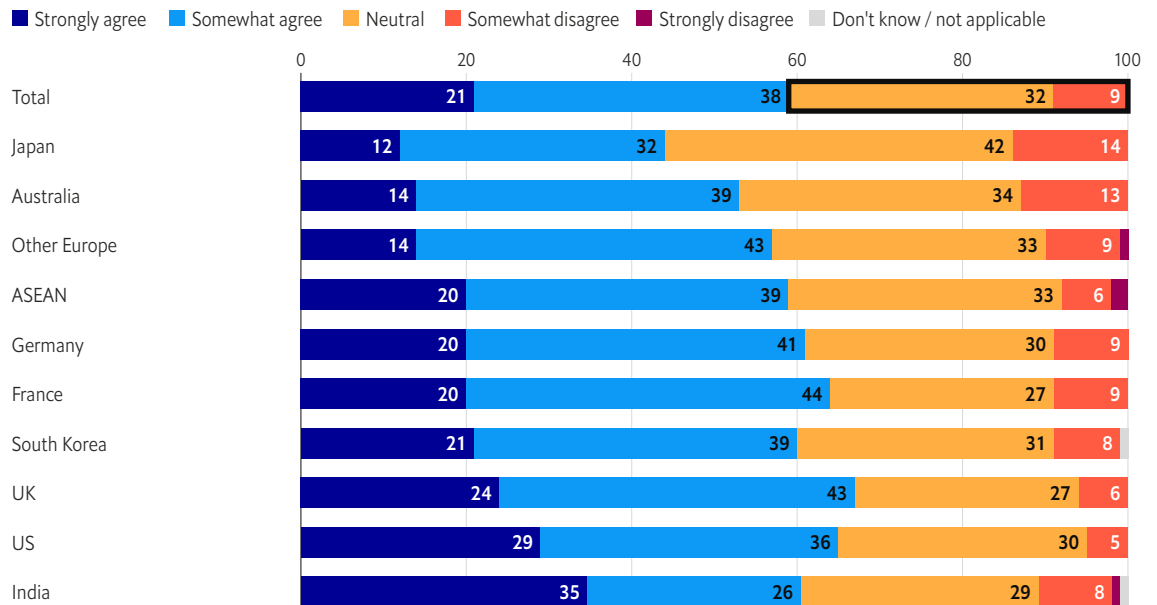
Centres of excellence: ‘forcing functions’ for co-ordination

While technical innovations are part of the governance toolkit, companies also need to consider forming new internal structures and institutions to co-ordinate AI experiments and embed best practices. Many of the large multinational enterprise executives interviewed for this report highlighted the importance of unified systems, clear processes and strong governance.

Our survey found lagging progress, however: 40% of respondents acknowledged that their organisation’s processes to ensure AI safety and compliance are insufficient (see figure 13). Companies of all sizes report similar perceptions about the quality of their governance. However, differences across geographies were more pronounced, with respondents in India (35%) and the US (29%) most likely to strongly agree that their processes were sufficient compared with the global average (21%).

Figure 13: Four in ten say their organisation’s AI governance is insufficient

My organisation has implemented sufficient processes to ensure AI safety and compliance. [Agree or disagree]
Percentage of all respondents



Source: Economist Impact
ASEAN includes Malaysia, Philippines, Singapore and Thailand
Other Europe includes Denmark, Finland, Italy, Netherlands, Norway, Spain and Sweden

“With GenAI, there is an opportunity to revisit an entire business process and reimagine what it could be. It encourages leaders to challenge past decisions and accept that there may be a newer, better way.”

Ryan Snyder, senior vice president and chief information officer, Thermo Fisher Scientific

Centres of excellence (COEs) are emerging as a popular approach to achieving systematic oversight. COEs have become key enablers of AI-driven insights and use cases within organisations. Beyond acting as the central control unit of AI, they monitor, implement and share best practices for implementation across departments and business units. They also foster exploration and innovation for building and operationalising AI.³²

Accenture’s Mr Ramani relates that his company collaborated with one national bank to build a COE to manage guidelines and protocols on how AI is used and adopted across the bank’s business units. While centres can forge consistency in practices and create guidelines and guardrails, they are not just brakes on AI. They can actually be a kind of ‘forcing function’ for AI delivery across the enterprise, according to Mr Ramani, bringing together departments, from HR and legal to tech teams, to unlock potential. A COE is a collaborative effort bringing together multidisciplinary teams and backgrounds for a common goal. For example, Shell’s Analytics COE brings together multi-faceted teams and supports a community of hundreds of data scientists and thousands of what the company calls ‘AI enthusiasts’ coming together.

Technically, COEs develop, ‘productionise’ and monitor AI models with a lens on business impact and returns. They oversee core operations, including for data and ML, and seek architecture and infrastructure optimisations for operational efficiencies. Most importantly, they unify the organisation’s AI vision and ambitions across business units and channels, while ensuring that safety and standards are met.

Measures like COEs are examples of how companies can take the opportunity of the AI moment to revisit all aspects of governance and decision-making. “With GenAI, there is an opportunity to revisit an entire business process and reimagine what it could be,” says Mr Snyder at Thermo Fisher Scientific. “It encourages leaders to challenge past decisions and accept that there may be a newer, better way.”

32 Deloitte, “Is your AI center of excellence still a center of experimentation”, <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/consulting/us-is-your-ai-center-of-excellence-still-center-of-experimentation.pdf>

Calibrating humans and machines

The march of the machines has long stirred fears of widespread job losses as automation threatens to render human workers obsolete. Yet, as AI technologies permeate industries, the reality is more nuanced. Many business leaders emphasise that AI adoption will not lead to layoffs. “AI means more efficiency in our business, which means we can create more stories—not by using AI to generate content but by automating mundane or manual tasks,” says Mr Bhakta at Condé Nast. “So instead of ten a week, we could produce 30 a week, which is actually good for the business. We’re not looking to shrink, we’re looking to grow. Doing more with the number of people we have is the best outcome.”

As mass AI adoption takes hold, there will be more demand for specialised roles like data scientists, developers and MLOps experts. Companies are already paying premiums for workers with AI-related abilities, and those who acquire AI skills see wages rise by an average of 21%.^{33,34} One 2030 forecast estimates that 375 million workers may need to switch occupations as their line of work is disrupted by digitalisation, automation and AI.³⁵


Only one in six firms are confident in their ability to attract the expertise needed to take full advantage of AI technologies. Those who succeed are deploying AI at a much faster rate.

Securing talent is one of the greatest challenges for firms deploying AI, our survey found; only one in six firms are confident in their ability to attract the expertise needed to take full advantage of AI technologies. Those who succeed are deploying AI at a much faster rate.

Meanwhile, the proliferation of GenAI is also seen as a democratising force. More than 97% of enterprise architects in our survey predict that non-technical staff will use natural language programmes to interact with complex datasets within the next three years, and 58% say that natural language will be the primary—or only—way they do this. Unilever, for instance, is investing in both people and culture through a programme that has trained almost 20,000 people in the use of GenAI, a vital step for a company operating in a dynamic sector where demand for new skills and competencies is rising, according to Mr Hill.

“A data and AI culture helps all parts of the business understand that we prioritise data-driven decisions, and that’s what will help us gain the insights that will improve performance,” says Mr Roberts at Seven West Media. “[Our] priority is to train and upskill our teams and talent around GenAI, so that it becomes a process and a product that’s internalised.”

- 33 Foote Partners LLC, “2024 IT Skills & Certifications Pay Index”, <https://footepartners.com/collections/pay-index>
- 34 Fabian Stephany, Ole Teutloff, “What is the price of a skill? The value of complementarity”, *Research Policy*, Volume 53, Issue 1, 2024, <https://doi.org/10.1016/j.respol.2023.104898>.
- 35 McKinsey, “Retraining and reskilling workers in the age of automation”, 2018, <https://www.mckinsey.com/featured-insights/future-of-work/retraining-and-reskilling-workers-in-the-age-of-automation>



Firms like Condé Nast and Frontier are prioritising engagement with employees, emphasising the advantages of automation rather than presenting it as a threat. Their approach highlights how AI can ease workloads and enhance job performance rather than eliminate positions.

Some firms are embedding AI tools into office software to ease workloads. “Our North Star is to democratise data. Everyone at the company should be able to leverage data that drives decisions ... without being a data engineer or a data scientist,” says Mr Bensaïd at Rivian. “Traditional AI involved specialised teams in expert environments,” says Mr Bucher at Novartis. “GenAI is accessible to many users, with self-explanatory tools and understandable use cases.”

Once a new technology has been developed and proven to be successful, the crucial step for a company is to develop self-service tools so employees can benefit, says Mr Casado at Repsol. “Technicians will deal with the technology—the real challenge is making it easy for the people to use [it],” he says.

The Repsol Artificial Intelligence Products hands employees the keys to AI, letting them tackle complex data tasks solo, without technical knowledge.³⁶ Staff can analyse performance and optimise operations and organisational change with natural language processors. These tools promise autonomy and agility, allowing Repsol to streamline operations while turning its workforce into “citizen data scientists”.

Self-service platforms can empower workers to handle tasks once left to specialists, like processing payroll, scheduling meetings or crunching numbers. The appeal for employers is more efficiency, fewer middlemen and lower costs. However, simply making AI available is not enough. Broader use of AI requires broader data literacy. The biggest data management challenges identified by architects in our survey, for instance, include controlling data access (49%) and protecting sensitive data (42%).

“Our North Star is to democratise data. Everyone at the company should be able to leverage data that drives decisions ... without being a data engineer or a data scientist.”

Wassym Bensaïd, chief software officer, Rivian



³⁶ Repsol, “The RAIP Project develops Repsol’s own AI solutions with company-wide impact and value”, <https://www.repsol.com/en/technology-and-digitalization/digital-transformation/digital-program/raip/index.cshml>

“We’re moving away from data literacy being a specialised skill,” says Mr Garciga at the US Army. All employees must understand data protection, platform use and centralised data storage. The US Army has incorporated these into its leadership training and has recognised that managing sensitive data is crucial to the Army’s AI use cases. One such use case involves reviewing documents to be declassified under the Freedom of Information Act, which allows for important documents related to events like the war in Afghanistan to be released to the public. “AI can improve the accuracy of this process and help us cover more documents, increasing transparency,” says Mr Garciga.

“We’re moving away from data literacy being a specialised skill. All employees must understand data protection, platform use and centralised data storage. The US Army has incorporated these into its leadership training and has recognised that managing sensitive data is crucial to the Army’s AI use cases.”

Leonel Garciga, chief information officer, US Army

Human judgement will remain critical for monitoring AI’s outputs, and companies must calibrate the right mix of people and machines. Two-thirds of organisations in our survey say they are actively experimenting to find the balance between humans and AI, and 80% say these efforts could be stronger. Mr Sato at NTT Docomo emphasises that “individuals using GenAI must have the ability to independently evaluate its outputs and handle them responsibly.”

Progress is most advanced in the financial services industry, where AI must fall within strict regulatory guidelines, our survey found. “As organisations begin using AI to inform key decisions, they must still work within agreed upon risk models. For example, any approach that TD takes to leverage AI must be within our risk appetite and regulatory obligations,” says Mr Martin at TD. That holds true for customer-facing interactions too. “In the old world, trust was established by walking in the branch, getting to know the manager and talking about the product,” says Mr Kapoor at Mahindra Group. “Now trust has to be established through cybersecurity and data privacy, and the customer needs to know that data will not be misused for any purpose.”

In sensitive sectors with high-risk profiles and low risk appetites, computer-generated errors will often provoke strong negative perceptions. As noted by Flo Health’s Mr Bugaev, medical specialists can also make mistakes or come to different answers to the same question, much like AI models might. But tolerance for errors and disparities in AI is understandably lower compared with tolerance for the same in medical professionals.

“Individuals using GenAI must have the ability to independently evaluate its outputs and handle them responsibly.”

Takaaki Sato, senior executive vice president, chief technology officer, NTT Docomo



Key takeaways

Data innovations enhance AI governance. Techniques like synthetic data and RAG help organisations address data limitations, reduce bias and improve AI system quality. These techniques can enhance performance, preserve privacy and enrich AI models with real-time information.

Centres of Excellence (COEs) drive systematic AI oversight. COEs act as central control units for AI implementation, monitoring best practices and fostering innovation across departments. They unify an organisation's AI vision, ensure safety standards are met and optimise core operations like MLOps. COEs bring together multidisciplinary teams, encouraging collaboration and reimagining business processes.

Balancing human expertise with AI capabilities is crucial. Organisations are focusing on upskilling employees and creating self-service AI tools to democratise data usage. While AI adoption increases efficiency, human judgement remains critical for monitoring AI outputs and handling sensitive data. Companies must calibrate the right mix of human and machine intelligence, especially in highly regulated industries where trust and risk management are paramount.

Conclusion

Internal use cases and pilots are the necessary testing ground. While there is an understandable desire to deploy AI at scale, companies are conscious of risks and unknowns; experts advocate internal use cases initially, through sandboxes and iteration. This learning phase also allows them to decide which mix of models and tools works best for which ends. Pilot selection should be determined by risk thresholds, to solve clear problems and within well-defined scopes.

Fixing the plumbing. To build a robust data and technical foundation, companies need to invest in an infrastructure refit to support high-volume, varied data. This will stress the rigid data storage systems of yesteryear, which cannot manage the diversity of data, real-time usage and wider utilisation in this more democratised era. Larger companies and conglomerates especially must decide the right balance between centralisation and devolution as they look to harness cloud platforms.

AI needs KPIs. To avoid profligate spending and the inevitable trough of disillusionment that could follow, companies need to develop intelligent ROI metrics and a disciplined approach to spending. Productivity, cost control, revenue per worker, user experience and reduced staff burnout are all meaningful metrics. The most successful will use AI to drive new revenue streams and business model innovation.

Use flexible timeframes for achieving returns. KPIs and metrics matter, but companies should not necessarily seek quick returns. Experiments involve wrong turns and tinkering. Companies in heavily regulated sectors like healthcare have to operate within structured systems of regulatory oversight. It takes time to develop the high-quality data sets and governance needed to power effective AI. There is also merit in 'watchful waiting' to see how the AI product and service ecosystem evolves.


Use a mix of models to sharpen your edge: in-house data and organisational know-how. As foundation models become commoditised, usage will be a baseline, not a boost. All organisations need to find their edge—their unique recipe of data and know-how. In highly specialised industries, off-the-shelf models will fall short of quality and compliance requirements. Companies can mix and match open- and closed-source AI, and private or public cloud, based on performance, governance, cost and capacity, to find the right blend.

Establish robust governance frameworks. Organisations need clear guidelines, protocols and oversight mechanisms, and institutional innovations like Centres of Excellence (COEs), to manage AI deployment and mitigate potential risks. COEs can also be a forcing function for improving co-ordination and coherence overall. Technical innovations like synthetic data and RAG can also strengthen model performance. It is crucial to recognise the importance of human oversight of AI systems, especially in high-stakes domains.

Calibrate human versus machine. Engaging with employees in the design and deployment of AI systems is critical to secure buy-in and adoption. That means communicating how AI can help workers rather than carry out their work for them. Identifying AI champions can also help build momentum. New skills and competencies can be developed through bite-sized programmes and exercises, with companies bridging the gap between AI specialists and the workforce. AI should also be integrated into leadership programmes. Last, but not least, organisations should encourage a critical outlook to avoid automation bias and keep human judgement at the fore.



While every effort has been taken to verify the accuracy of this information, Economist Impact cannot accept any responsibility or liability for reliance by any person on this report or any of the information, opinions or conclusions set out in this report. The findings and views expressed in the report do not necessarily reflect the views of the sponsor.



LONDON

The Adelphi
1-11 John Adam Street
London WC2N 6HT
United Kingdom
Tel: (44) 20 7830 7000
Email: london@economist.com

GENEVA

Rue de l'Athénée 32
1206 Geneva
Switzerland
Tel: (41) 22 566 2470
Fax: (41) 22 346 93 47
Email: geneva@economist.com

SINGAPORE

8 Cross Street
#23-01 Manulife Tower
Singapore
048424
Tel: (65) 6534 5177
Fax: (65) 6534 5077
Email: asia@economist.com

NEW YORK

900 Third Avenue
16th Floor
New York, NY 10022
United States
Tel: (1.212) 554 0600
Fax: (1.212) 586 1181/2
Email: americas@economist.com

DUBAI

Office 1301a
Aurora Tower
Dubai Media City
Dubai
Tel: (971) 4 433 4202
Fax: (971) 4 438 0224
Email: dubai@economist.com

WASHINGTON DC

1920 L street NW Suite 500
Washington DC
20002
Email: americas@economist.com

HONG KONG

1301
12 Taikoo Wan Road
Taikoo Shing
Hong Kong
Tel: (852) 2585 3888
Fax: (852) 2802 7638
Email: asia@economist.com