

Generative AI and Creating the Art of the Possible

Implementing gen AI today while planning for tomorrow

Table of Contents

About the Author	3
Introduction	3
Overview	3
Modern technology strategy	6
People and Culture	7
Risks, ethics, and regulation	10
Innovation and reinvention	14
Conclusion	18
Final Takeaway	19
Sources	20
Resources	20



About the Author

Vivek Bhalla is the lead for MongoDB's core database product marketing team. He previously spent three years building the company's Competitive Intelligence team.

Prior to MongoDB, Vivek worked in technical support, product management, and director roles at companies like Micromuse, IBM, and AIOps pioneer Moogsoft. He was also a Research Director at Gartner, during which time he was lead author for published research including Magic Quadrants, Hype Cycles, and Market Guides. Covering the IT Operations space, and being the global lead for Network Management, he defined concepts like [Network Automation](#) (NA) and [Unified Communication Monitoring](#) (UCM), and he helped coin the term '[AIOps](#)'—Artificial Intelligence for IT Operations.

Introduction

Over the past 25 years, the core elements of any transformative change have been people, culture, process, and technology. Generative AI (gen AI) follows this pattern, but stands out as potentially more significant than any previous technological transformation. Put simply, gen AI has the potential to impact every industry and to fundamentally alter how we live, work, and communicate.

Readers can use this white paper to:

- Determine the challenges identified by industry and AI experts
- Understand how executives and IT leaders tasked with addressing those challenges are developing their approaches and strategy
- Formulate their own best practices concerning AI, data management, and application development

Gen AI's rapid advancement is causing disruptions in traditional job roles, requiring people across industries to adapt to new tasks. In this ever-evolving environment, there are endless possibilities stretching before us. The path ahead may be uncertain, but one thing remains clear: the journey is just beginning.

Overview

With industry interest in gen AI rapidly evolving and the diversity of technology proliferation in abundance, it can be challenging to discern hype from real-world use cases that deliver actual value to businesses.


Despite a decline in overall AI private investment last year, funding for gen AI surged, nearly octupling from 2022 to reach \$25.2 billion. Major players in the gen AI space—including OpenAI, Anthropic, Hugging Face, and Inflection—reported substantial fundraising rounds.¹

¹Source: Stanford Institute for Human-Centered Artificial Intelligence (HAI)



Private investment in generative AI, 2019–23

Source: Quid, 2023 | Chart: 2024 AI Index report




Since 2011, the number of AI-related projects on GitHub has seen a consistent increase, growing from 845 in 2011 to approximately 1.8 million in 2023. Notably, there was a sharp (+59.3%) rise in the total number of GitHub AI projects in 2023

alone. The total number of stars for AI-related projects on GitHub also significantly increased in 2023, more than tripling from 4.0 million in 2022 to 12.2 million.

Number of GitHub AI projects, 2011–23

Source: GitHub, 2023 | Chart: 2024 AI Index report




A 2023 McKinsey survey reveals that 42% of surveyed organizations report cost reductions from implementing AI (including gen AI), and 59% report revenue increases. Compared to the

previous year, there was a 10 percentage point increase in respondents reporting decreased costs, suggesting AI is driving significant business efficiency gains.



Cost decrease and revenue increase from AI adoption by function, 2022

Source: McKinsey & Company Survey, 2023 | Chart: 2024 AI Index report



In June 2024, KPMG polled 100 US-based C-suite and business leaders, representing organizations with an annual revenue of \$1 billion or more. 80% of respondents said they see gen AI spending as key to gaining a competitive advantage and boosting market share.

While AI has already demonstrated its potential in automating tasks and enhancing decision-making processes, the integration of gen AI has the potential to take innovation to the next level. This raises the question: how can organizations bridge the gap between the theoretical promise of gen AI and its tangible impact to make informed decisions and to develop the strategies that will harness the full promise of this groundbreaking technology?

The modern C-level executive has a fundamental role in both the immediate and ongoing adoption and implementation of gen AI. As such, it is vital to set one's organization up for transformative innovation to take advantage of all that gen AI

offers, yet maintain awareness and agility to avoid the perils and risks of being sucked into hyperbole that is also associated with such a disruptive technology.

IT leaders face a multitude of major hurdles to the effective adoption and scaling of gen AI, including a talent shortage/expertise, poor data quality or integration, lack of comprehensive AI governance, and risk mitigation and control. Escalating costs and unclear business value are other leading causes of gen AI project failure, according to Gartner. So while cost reductions are possible, without correct oversight to drive the right goals and relationship with gen AI from the outset, the implementation is prone to failure.

For technology CXOs, step changes pose as many business challenges as they do opportunities. If organizations are to benefit from AI, then they need the infrastructure fit for purpose, a new culture, governance—especially around data—and a new relationship with technology.



Modern technology strategy

Enterprises are rapidly adopting gen AI to increase productivity and efficiency, however many are not taking a strategic approach to implementing the technology. As a consequence, many projects fail or end up costing far more than they should, without a return on investment (ROI).

Through 2025, Gartner predicts at least 30% of gen AI projects will be abandoned after proof-of-concept due to issues like poor data quality, inadequate risk controls, escalating costs, or unclear business value.

Creating success means CXOs have to analyze and implement gen AI in ways that will allow it to positively supplement existing operations, or even create entire new lines of business. Gen AI will do two things to organizations:

1. Accelerate the automation of everyday tasks, cutting the amount of manual and repetitive—and typically not value-adding—activities that skilled team members do.
2. Create new business models. Gen AI is more than an automation tool, as has already been seen with its ability to diagnose illnesses from large data sets, to predict rogue waves that can sink ships, or to communicate on a human level.

These opportunities will only be achieved in organizations that have the proper infrastructure in place. A modern technology stack requires substantial computing resources, advanced data systems, essential development tools, and access to models, either open-source or via commercial APIs. These elements are crucial for building scalable, innovative solutions.

Many organizations have yet to determine which technology environment is best suited to run the different parts of a gen AI process and workload. Part of that challenge is that organizations have yet to decide which gen AI applications are most suited to their business and vertical market. This is understandable, as the pace of development is rapid, and the number of solutions available is extensive. With increasingly advanced models being released all of the time, the situation is exacerbating.

Just as with mobile and cloud computing before it, gen AI will trigger a wave of technology infrastructure modernization. For organizations to extract business value from gen AI, they will require an interconnected data environment. Therefore, it is no surprise to learn that more than half of CXOs say they need to improve their data transfer abilities between multi-cloud, data center and edge computing environments.² The same study finds that most CXOs cannot pinpoint the infrastructure modernization plan needed to support AI workloads.

Investment in and modernization of technology infrastructure is going to become a continual and long-term program for technology leaders in order to meet business goals and customer expectations. This gold rush towards AI-optimized organizations will not forego the need for technology CXOs to be mindful of budgets and risks.

The demands of the market are brutal; macroeconomic conditions, the rapid pace of competition and innovation, and new technology shifts like AI are putting immense pressure on business and technology leaders. They feel tasked with an impossible balance: how to invest for the future, for innovative new experiences while finding ways to keep the lights on and continue to refine towards operational excellence, to streamline the business, mitigate risks, and reduce costs?

Identifying those technology partners with a track record of enabling modernization will be key. CXO and IT leaders responsible for this endeavor must:

- Shift underlying infrastructure to the cloud
- Migrate to modern data platforms that unlock previously untapped productivity
- Adopt new architectures and platforms by moving to microservices patterns
- Align behind new IT models and processes, such as Agile and DevOps

²Source: [Comparethecloud.net](https://www.comparethecloud.net)



People & culture

“On every dollar of technology we need to invest three to five in human beings”

— Lareina Yee McKinsey, Senior Partner, Chair of the McKinsey Technology Council—McKinsey & Company, August 2024

“Technology is the easy bit” is a frequent quip of CXOs, the greatest hurdle for business technology leaders and their organizations is the culture change—the same will be true of gen AI. The conversational and broadly accessible capabilities of gen AI make it both an easy tool for end-users to engage with, but also an increased security risk.


In a January 2024 Gartner poll, nearly two-thirds of organizations were reported to be using gen AI across multiple business units, a 19% increase from September 2023. 40% of organizations had deployed gen AI in more than three business units. The same poll also cited Customer service (16%), Marketing (14%) and Sales (12%) as the primary business functions that have adopted or intended to invest in gen AI in some form. These kinds of non-technical employees can be inadvertently uploading intellectual property or customer data with ease, instantly exposing the business to regulatory and market risk.

Such an intuitive engagement model heralds a new relationship that business end-users have with technology. For example, Gartner predicts that by 2028, one-third of interactions with gen AI services will invoke action models and autonomous agents for task completion, and Thomson Reuters anticipates that all professionals will have a gen AI assistant within five years (2029). One needs to consider gen AI not as a technology to be implemented in the traditional sense but as a coworker that, just as with any colleague, has to be interviewed and coached on the culture of their team and organization.

Cultural change is not limited to business end-users either. We do not know if assisted engineering will replace engineers. However, it is evident that engineers who familiarize themselves and embrace AI tools (e.g. Copilot) will supersede those who do not.

Cross-study comparison of task completion speed of Copilot users

Source: Cambon et al., 2023 | Chart: 2024 AI Index report



In 2023, several studies assessed AI's impact on labor, suggesting that AI enables workers to complete tasks more quickly and to improve the quality of their output.³ These studies also demonstrated AI's potential to bridge the skill gap between low- and high-skilled workers. Still other studies caution that using AI without proper oversight can lead to diminished performance.


When people ask about talent today, they are often concerned with how roles are changing—whether certain skill sets are in ascendancy or decline. Technology does change the job market, in doing so it also creates new opportunities and new roles.

Companies must focus on finding and developing the right skills to build and grow these new technologies. At the same time, organizations must consider how these changes will affect the workforce holistically. Understanding these shifts is critical to making the most of new technology in the workforce.

It should be noted that while AI has surpassed human performance on several benchmarks—including image classification, visual reasoning, and English language understanding—it trails behind on more complex tasks. These include competition-level mathematics, visual commonsense reasoning, and planning.

Select AI Index technical performance benchmarks vs. human performance

Source: AI Index, 2024 | Chart: 2024 AI Index report



A survey from Ipsos shows that, over the last year, the proportion of those who think AI will dramatically affect their lives in the next three to five years has increased from 60% to 66%. Moreover, 52% express nervousness toward AI products and services, marking a 13% point rise from 2022.

In America, Pew data suggests that 52% of Americans report feeling more concerned than excited about AI, rising from 38% in 2022.


Ipsos survey results also indicate only 37% of respondents feel AI will improve their job. Only 34% anticipate AI will boost the economy, and 32% believe it will enhance the job market.

³Source: Stanford Institute for Human-Centered Artificial Intelligence (HAI)



Global opinions on the impact of AI on current jobs, 2023

Source: Ipsos, 2023 | Chart: 2024 AI Index report



Gen AI will be a step change in enterprise IT. It will, therefore, also be a step change in technology leadership. To achieve success, organizations will need the right infrastructure, governance, and culture in place. These three require an understanding that gen AI is a business issue and not one solely for IT to resolve.

CXO and IT leaders contending with how to introduce gen AI to their organizations should:

- Designate an AI lead and liaison in each department to translate business need to technical capabilities
- Determine those technology and service partners that have existing *production* deployment experience of gen AI adoption through multiple customer engagements

- Formulate best practices based on prior lessons learned from the path tread by others rather than forge one's own path without any 'guardrails'
- Instigate company-wide training on prompt engineering, for all employees irrespective of seniority, job title, or role
- Employ end-user feedback loops to ensure domain-specific use case knowledge and sentiment is captured to help fine tune and improve model accuracy and relevance




Risks, ethics, and regulation

Many leaders acknowledge that risks exist with gen AI, yet often struggle to define or articulate them clearly. Among the most pressing concerns are inaccuracy, intellectual property infringement, cybersecurity, individual privacy, regulatory compliance, explainability, fairness, and bias amplification. This highlights a crucial need for better awareness and communication around these risks.

Global surveys conducted by Stanford University, Thomson Reuters, and Forrester cite privacy, security, and governance as the greatest challenges organizations consider themselves faced with when utilizing AI. In a Forrester survey comprising 3,474 business and technology decision-makers, privacy was rated above cost, a lack of skills, and several other concerns. Thus far, most companies have only mitigated a portion of this exposure.

Relevance of selected responsible AI risks for organizations by region

Source: Global State of Responsible AI report, 2024 | Chart: 2024 AI Index report



Risk mitigation constitutes a multitude of considerations, such as the specific use case and the consequence on individuals of any outcome that gen AI helps an organization come to. Gen AI can harm individuals and communities if adequate data protection is not implemented.⁴ Therefore, organizations must define risk impact by tying the

consequence of privacy protection deficiencies to those of application safety deficiencies. Determining where personal and sensitive data is involved is a prerequisite. This raises considerations across three elements of potential risk exposure:

Risk exposure	Details
Training models	Large language models (LLMs) scrape copious gigabytes of data often across a variety of internet sources. This raises the risk that personally identifiable information (PII) data is inadvertently incorporated into these models. ⁵ Such occurrences are liable to fall foul of legal safeguards. This risk is compounded when the training set data or its sources are opaque and obfuscated by providers.
Internal data sets	Organizations employing their own data sets to train a gen AI app must still identify and classify their own sensitive information. Intellectual property, confidential corporate details, or employee PII must all be treated differently. Governance of this can be easily overlooked and difficult to maintain and track.
Output	The output of a gen AI app may contain new content that, with inference, may glean sensitive or personal insights. An EU court has already set precedent that organizations are liable for the sensitivity of any newly generated data.

⁴Source: [EPIC](#), May 2023 ⁵Source: [Washington Post](#), April 2023




The Foundation Model Transparency Index shows that AI developers lack transparency, especially regarding the disclosure of training data and

methodologies. This lack of openness hinders efforts to further understand the robustness and safety of AI systems.


Foundation model transparency total scores of open vs. closed developers, 2023

Source: 2023 Foundation Model Transparency Index



In research conducted by Dimensional Research in 2024, where over 500 executives and AI professionals were surveyed across diverse sectors, 95% of companies using gen AI for business applications felt greater security measures were required to protect their data. The same research indicated that 94% of respondents felt data protection was a primary consideration where utilization of gen AI was concerned. With the potential for sensitive data being exposed, data control and access is cited as a prerequisite.

In your experience, what types of security-related controls are needed for generative AI to be used in business applications?



KEY TAKEAWAY

Data control and access lead security controls needed for generative AI

In tandem with the greater risk, global AI-related regulation and governance has risen significantly in the past year and over the last five years. In July 2024, the full and final text of the EU AI Act, the European Union’s landmark risk-based regulation for applications of AI, was published in the bloc’s [Official Journal](#). This followed the establishment of the EU’s [AI Office](#) in June 2024, the ecosystem-building and oversight body providing regulatory framework. And in November 2023, the UK government published [The Bletchley Declaration](#), encapsulating the aims of 29 countries, including China, USA, Japan, Korea, India, Singapore, Israel and EU member states, that attended an AI Safety Summit, the goal of which was to reach global consensus on how to tackle the risks that AI poses both now and in the future.


What’s more, in 2023, there were 25 AI-related US regulations, up from just one in 2016. In 2023 alone, the total number of AI-related US regulations grew by 56.3%⁶. Ensuring adherence to these regulations will place an overhead on organizations seeking to adopt gen AI.

⁶Source: Stanford Institute for Human-Centered Artificial Intelligence (HAI)



Number of AI-related regulations in the United States, 2016–23

Source: AI Index, 2024 | Chart: 2024 AI Index report



Multiple researchers have shown that the generative outputs of popular LLMs may contain copyrighted material, such as excerpts from The New York Times or scenes from movies. Gartner predicts that by 2026, 75% of businesses will use gen AI to create synthetic customer data, up from less than 5% in 2023. Whether such output constitutes copyright violations is a pressing legal

question, and is compounded by gen AI models that “[hallucinate](#)”. Gen AI hallucinations are scenarios in which models produce inaccurate or fictional information. Hallucinations can manifest as fabricated events, misleading information, imaginary scientific concepts, or incorrect quotes to real or fictional characters.

Identical generation of a movie character

Source: [Marcus and Southen, 2024](#)



During a MongoDB gen AI Executive CXO Briefing held in 2024, CXOs told us that over the next two years, data modeling, and data security governance challenges will be high on their agenda. This tallies with the findings of the Digital Leadership Report by Nash Squared, which states that 36% of CXOs are concerned about data privacy being compromised by gen AI implementations. Concerns regarding hallucinations in the data was another commonly raised issue when MongoDB spoke with CXO leaders. The major concern being that this could impact customers and, therefore, the brand value of the business.

“Through 2026, one-third of enterprises will realize that a lack of AI and ML governance has resulted in biased and ethically questionable decisions.”

**— David Menninger, Executive Director,
Technology Research—Ventana Research,
January 2024**

Given the potential ramifications and liability this places on any organization, CXO and IT Leaders must strive to:

- Identify *trusted* technology partners that have not only validated their own technology but vetted *their* alliance partners to increase assurance and reduce risk
- Build a governance framework for gen AI that includes policies and training to expedite employee awareness, education, and confidence
- Alert against high-risk use cases, particularly where training models are liable to evolve rapidly and/or source data from dubious or ambiguous *sources*
- Promote data transparency and be clear about any retention periods
- Conduct extended risk assessments that are acutely mindful of the consequences of unintended model output, or the incorporation of sensitive data that may risk discrimination against individuals



Note: MongoDB is a founding member of the National Institute of Standards and Technology (NIST)’s AI Safety Institute Consortium (AISIC). This is a working group to support the development and deployment of safe and trustworthy AI. By identifying best practices to mitigate against the potential dangers of AI the AISIC looks to underpin future standards and policies.

The Consortium brings together over 200 AI creators and users, academics, government and industry researchers, and civil society organizations. The complete list of current Consortium members can be found on the NIST website.

⁷Source: [Ventana Research](#), January 2024



Innovation and reinvention

Senior leaders tend to overestimate the value of technology, including gen AI, in the short term yet underestimate its longer term potential. How does one strike the right balance?

Avoid repeating the mistakes of the past.


Top-down mandates from leadership on the percentage of apps that needed to be moved to the cloud resulted in poor outcomes. This was because the suitability of applications and workloads for the cloud were not fully considered, resulting in simple “lifts and shifts” that drove up business cost and risk. In the case of gen AI, these risks can be significantly elevated. Therefore, it is important to objectively assess where gen AI makes sense, and not pursue arbitrary metrics.

Gen AI is very different from traditional “predictive AI”. Because of these differences—for example probabilistic vs deterministic models—a number of CxOs MongoDB has spoken to agree with the premise that data scientists should not be involved with gen AI projects. Instead, build pods staffed with a business owner, product manager, and developers who own the gen AI “products” along with the controls necessary to ensure trust and quality.

In Dimensional Research’s survey it was revealed that as a consequence of security and data management concerns, 78% of surveyed companies had already deployed, or had planned to opt for, an in-house gen AI solution. It is a natural assumption that one will have more control over their own data if they have full ownership of the security measures taken. Yet organizations are typically unprepared for the scale and cost of adopting an in-house approach. This will have an impact on one’s ability to maintain a pace

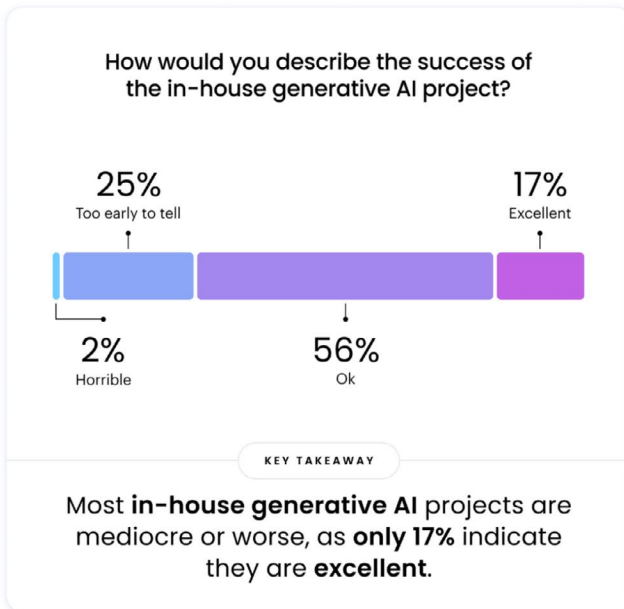
of innovation. Underestimated timescales and unexpected errors and technical challenges will also have a compound consequence.

Dimensional Research’s survey indicated that 61% of companies that opted for an in-house strategy have experienced inaccuracies. The impact being that expected value is not realized, nor are intended goals met. The result being wasted time, resources and business opportunity. Concurrently, competitors that have turned to previously field-tested commercial solutions, realize the intended value and deliver customer satisfaction.



Furthermore, only 17% of companies surveyed indicated that their in-house gen AI projects were “excellent”. Given the costs involved, attaining average or perfunctory levels of success, or worse, can only be deemed poor ROI.





A major contributory factor to the issue of accuracy was that only 10% of those surveyed reported a high confidence in building a quality in-house solution. Given the specialized skill sets and resources required, any organization treading this path can be overwhelmed by the complexity and pace of change involved with gen AI.


MongoDB's discussions with leading multinational professional services and consultancy executives revealed predicted costs upwards of \$1B USD

to train LLMs going forwards. These IT leaders anticipated only a few companies globally will feasibly have the capacity to perform this role.

According to AI Index estimates, the training costs of state-of-the-art AI models have reached unprecedented levels. For example, OpenAI's GPT-4 used an estimated \$78 million worth of compute to train, while Google's Gemini Ultra cost \$191 million for compute.

Estimated training cost of select AI models, 2017–23

Source: Epoch, 2023 | Chart: 2024 AI Index report




An additional consideration is that on 10 select AI benchmarks, closed models outperformed open ones, with a median performance advantage of 24.2%. Differences in the performance of closed

and open models carry important implications for AI policy debates as well as technology adoption for individual organizations.

Performance of top closed vs. open models on select benchmarks

Source: AI Index, 2024 | Chart: 2024 AI Index report




In 2023, a total of 149 foundation models were released, more than double the amount released in 2022. Of these newly released models, 65.7% were open source, compared to only 44.4% in 2022 and 33.3% in 2021. However, in terms of

performance, closed models outperformed open ones. Tracking and maintaining awareness of this proliferation and the dynamics of the market is an onerous overhead for a single organization to take on.

Foundation models by access type, 2019–23

Source: Bommasani et al., 2023 | Chart: 2024 AI Index report



Given the challenges involved, organizations must carefully weigh the benefits and limitations of opting for an in-house solution versus the advantages posed by commercial offerings.

Yet, selection of the *right* commercial solution is also key. Research from the AI Index reveals a

significant lack of standardization in responsible AI reporting. Leading providers, including OpenAI, Google, and Anthropic, primarily test their models against different responsible AI benchmarks. This practice complicates efforts to systematically compare the risks and limitations of top AI models.

Reported responsible AI benchmarks for popular foundation models

Source: AI Index, 2024 | Table: 2024 AI Index report

Responsible AI benchmarks	GPT-4	Llama 2	Claude 2	Gemini	Mistral 7B
TruthfulQA	✓	✓	✓		
RealToxicityPrompts	✓			✓	
ToxiGen		✓			
BOLD		✓			
BBQ			✓	✓	

While the focus may be towards LLMs and foundation models, selection of the most suitable data management platform is also key. In Thomson Reuters’ Future of Professionals 2024 report, 60% of respondents cited an explosion in the volume of data that they encounter will have a high or transformational impact on their work, a jump of 9 percentage points compared to 2023.

Better management of data is a primary driver of many legacy modernization initiatives. Research from The McKinsey Global Institute concludes data driven organizations are 23 times more likely to acquire customers, six times as likely to retain customers, and 19 times as likely to be profitable as a result.

CXO and IT leaders must identify the right data management platform that:

- Delivers the flexibility that gen AI demands, while also ensuring the robustness and security that any enterprise-grade workload expects
- Rapidly integrates multiple disruptive technologies, such as AI, IoT, mobile and social channels, cloud platforms, across their entire enterprise estate
- Embraces digital business models to generate new revenue streams and remove the impediment and technical debt of legacy platforms and process

Conclusion

Senior executives are at a critical juncture as the gen AI wave sweeps the IT industry. This is an inflection point when pivotal decisions will be made to establish long-term trends and determine a company's trajectory.

“We’re at a time where every company is facing a crucible moment when it comes to AI. AI can completely reimagine how you run your business, how you engage with your customers, and how you become more nimble and efficient”⁸

— Dev Ittycheria, CEO—MongoDB, May 2024

Every day, something new regarding gen AI is in the news. A new startup gets funding, an organization is rolling out a new AI-based service or capability, or a new concern regarding safety, ethics or regulation is cited. As with any technology adoption, the potential to over-extrapolate in the short term yet underestimate the long term persists.

There is also the consideration that AI technology is expensive, in terms of both, capex and opex costs. However, with use cases still in their infancy, the ROI model remains unclear. Daron Acemoglu, Institute Professor at MIT, anticipates truly transformative changes progressing over a decade, with nearer term impact limited to pure mental tasks. Academic studies estimate cost savings from the use of AI ranging anywhere from 10% to 60%.⁹

Given this premise, it is easy to feel overwhelmed, particularly given the rate and pace of innovation that is occurring in this space coupled with the multifaceted considerations on several levels—technology, people, process, culture, governance, ethics and regulation. Concurrently, there is the risk of *not* acting. In a Thomson Reuters Future of Professionals Report, almost one third (31%) of professionals working in legal, tax and accounting, and risk and compliance industries felt their firm or department was moving too slowly in terms of AI adoption and feared this leading to a loss of competitive advantage. If one does not take advantage of new technology, the consequence could mean being marginalized or completely

disrupted as others steal an insurmountable competitive advantage.

How does one validate a reference architecture? Will such architectures cover the use cases their organization is looking to address? Are there suitable integrations with each technology partner involved? How extensive is the technical expertise and support to enable rapid implementation and upskill internal staff?

Identifying the right technology partners is more pervasive than selecting the technical solution that ticks the most features on a pick list today. When determining technology partners, organizations should ensure that they:

- Prioritize providers that think beyond the technology and encompass people, process *and* culture into their AI framework.
- Avoid being locked into a single vendor's ecosystem—flexibility and adaptability is key.
- Pinpoint agile technology companies with a track record of supporting enterprise-grade production deployments.
- Challenge technology companies to demonstrate an ability and willingness to collaborate with third parties, be they hyperscalers, other software vendors, systems integrators or value added resellers.
- Conduct independent research to validate initial findings. Use [industry reports](#), speak with analysts and [compare findings with peers](#).

⁸Source: [CRN](#), May 2024 ⁹Source: [Goldman Sachs](#), June 2024



Final takeaway

If the considerations outlined above are top of mind for you and your organization, MongoDB is at hand to discuss your specific requirements and provide guidance.

The MongoDB AI Applications Program (MAAP) was launched to provide strategic advisory, professional services, and an integrated end-to-end technology stack from MongoDB and the MAAP ecosystem of companies. It offers customers a breadth of resources to put AI applications into production and a unified support system to help customers quickly build and deploy AI applications. The MAAP ecosystem includes industry-leading consultancies, foundation model, cloud infrastructure, generative AI framework, and model hosting providers.

“Generative AI is a transformative technology that has the promise to change how organizations of all sizes across industries conduct business, optimize their operations, and interact with their customers. AWS and MongoDB have been working together to integrate services for several years, helping customers leverage the strength of both companies to help developers innovate more rapidly. The MongoDB AI Applications Program will help build on the impact we’ve already been helping customers achieve.”



— Chris Grusz, Managing Director
Technology Partnerships, AWS

“MongoDB and Google Cloud have helped joint customers enhance their businesses with AI for many years, and we’re excited to extend our collaboration with support for the MongoDB AI Applications Program. Together, we will provide enterprises with the technical resources, frameworks, and governance tools required to more effectively build and deploy generative AI applications on Google Cloud’s AI-optimized infrastructure.”



— Stephen Orban, VP of Migrations, ISVs, and
Marketplace, Google Cloud

“Many customers want to get up and running quickly with generative AI but are sometimes overwhelmed by the many choices of technologies and the expertise required to build modern applications. The new MongoDB AI Applications Program will help organizations reduce the complexity and time it takes to build applications enriched with generative AI by offering them integrated technology solutions and hands-on support. We are excited to be a launch partner for the program to help our joint customers with MongoDB take advantage of this game-changing technology.”



— Alvaro Celis, Vice President of Global ISV
Commercial Solutions, Microsoft

MAAP ECOSYSTEM

accenture

ANTHROPIC

anyscale

arcee.ai

cohere

CREDAL

Fireworks AI

gravity

LangChain

LlamaIndex

NOMIC

PeerIslands

pureinsights

together.ai



Sources

- [2023 OverStack Developer Survey](#), May 2023
- [A Privacy Primer On Generative AI Governance](#), Forrester, July 2023
- [Digital Leadership report](#), Nash Squared, 2023
- [2024 U.S. CEO Outlook Pulse Survey](#), KPMG, April 2024
- Tech Crunch, April 2024: [Intel and others commit to building open generative AI tools for the enterprise](#)
- Constellation Research, April 2024: [Anthropic CEO Amodei on where LLMs are headed, enterprise use cases, scaling](#)
- The New York Times, April 2024: [AI Start-Ups Face a Rough Financial Reality Check](#)
- IT Pro, April 2024: [Gemini Code Assist could be Google's secret weapon to challenge GitHub Copilot](#)
- Runtime, April 2024: [Who gets to decide what "safe" AI means?](#)
- BenchANT, April 2024: [To Benchmark Vector Databases or to Get Sued for Breaching a DeWitt clause?](#)
- BBC, May 2024: [Business locked in expensive AI 'arms race'](#)
- [The State of AI](#) report June 2024
- [Gen AI: Too Much Spend, Too Little Benefit?](#), Goldman Sachs, June 2024
- [Future of Professionals Report](#), Thomson Reuters, July 2024
- [Gartner Survey Reveals CFOs and CEOs Identify AI as the Technology with the Greatest Impact in the Next Three Years](#), Gartner, July 2024
- [KPMG GenAI Study: The Path to Sustainable Returns](#), KPMG, July 2024
- [Market Observations on AI Governance](#), Ventana Research, August 2024
- [AI hallucinations: What are they?](#), ITPro, August 2024
- [Generative AI in the enterprise 2024 survey](#), Dimensional Research, 2024
- [What Generative AI Means for Business](#), Gartner, 2024
- Stanford Institute for Human-Centered AI: [AI Index Annual Report 2024](#)

Resources

For more information, please visit mongodb.com or contact us at sales@mongodb.com.

Case Studies (mongodb.com/customers)

Presentations (mongodb.com/presentations)

Free Online Training (university.mongodb.com)

Webinars and Events (mongodb.com/events)

Documentation (docs.mongodb.com)

MongoDB Atlas database as a service for MongoDB (mongodb.com/cloud)

MongoDB Enterprise Download (mongodb.com/download)