


A playbook for crafting AI strategy



Preface

“A playbook for crafting AI strategy” is an MIT Technology Review Insights report sponsored by Boomi. To produce this report, MIT Technology Review Insights conducted a global survey of C-suite and senior data executives across countries and industries. The report also draws on in-depth interviews conducted with business leaders on data and AI.

Adam Green was the author of the report, Teresa Elsey was the editor, and Nicola Crepaldi was the publisher. The research is editorially independent, and the views expressed are those of MIT Technology Review Insights.

We would like to thank the following executives and experts for their time and insights:

Kevin Collins, Founder and Chief Executive Officer, Charli AI

Amy Machado, Senior Research Manager, IDC

Matt McLarty, Chief Technology Officer, Boomi

SP Singh, Senior Vice President and Global Head, Enterprise Application Integration and Services, Infosys

About the survey

The survey forming the basis of this report was conducted by MIT Technology Review Insights in March 2024. The survey sample consists of 205 executives and data and technology leaders. Eleven industries are represented: financial services, manufacturing, IT and telecommunications, consumer goods and retail, pharmaceutical and health care, government, travel and hospitality, professional services, energy and utilities, transport and logistics, and media and marketing.


Nearly all survey respondents (88%) come from the C-suite. These include chief executive officers (20%), chief information officers (18%), chief technology officers (19%), and chief data officers (15%). The respondents' organizations are headquartered in North America (31%); Europe, the Middle East, and Africa (25%); Asia-Pacific (26%); and Central and South America (18%). All respondents work at organizations with more than US \$500 million in global annual revenue, with 73% representing organizations generating more than US \$1 billion, and 34% more than US \$10 billion.

CONTENTS

01	Executive summary	4
02	Partnering for success	6
	Selecting a vendor	7
	Finance-friendly AI	7
03	Counting the cost	8
	Spending expectations	9
	Measuring return on investment	10
04	Building a data core	11
	Data management: Tips and tactics	11
	Reckoning with legacy infrastructure	13
	Data lineage and liquidity	13
	Metadata	13
05	Acceleration versus caution	14
	Hallucinations, errors, and bias	14
	Cyber risk	14
	Data privacy and protection	14
	Rising regulatory tide	15
	Compliance challenges	16
06	Conclusion	17

01

Executive summary



Giddy predictions about AI, from its contributions to economic growth to the onset of mass automation, are now as frequent as the release of powerful new generative AI models. The consultancy PwC, for example, predicts that AI could boost global gross domestic product (GDP) 14% by 2030, generating US \$15.7 trillion.¹

Forty percent of our mundane tasks could be automated by then, claim researchers at the University of Oxford, while Goldman Sachs forecasts US \$200 billion in AI investment by 2025.^{2,3} “No job, no function will remain untouched by AI,” says SP Singh, senior vice president and global head, enterprise application integration and services, at technology company Infosys.

While these prognostications may prove true, today’s businesses are finding major hurdles when they seek to graduate from pilots and experiments to enterprise-wide AI deployment. Just 5.4% of US businesses, for example, were using AI to produce a product or service in 2024.⁴

Moving from initial forays into AI use, such as code generation and customer service, to firm-wide integration depends on strategic and organizational transitions in infrastructure, data governance, and supplier ecosystems. As well, organizations must weigh uncertainties about developments in AI performance and how to measure return on investment.

If organizations seek to scale AI across the business in coming years, however, now is the time to act. This report explores the current state of enterprise

“No job, no function will remain untouched by AI.”

SP Singh, Senior Vice President and Global Head, Enterprise Application Integration and Services, Infosys

AI adoption and offers a playbook for crafting an AI strategy, helping business leaders bridge the chasm between ambition and execution. Key findings include the following:

AI ambitions are substantial, but few have scaled beyond pilots. Fully 95% of companies surveyed are already using AI and 99% expect to in the future. But few organizations have graduated beyond pilot projects: 76% have deployed AI in just one to three use cases. But because half of companies expect to fully deploy AI across all business functions within two years, this year is key to establishing foundations for enterprise-wide AI.

AI readiness spending is slated to rise significantly. Overall, AI spending in 2022 and 2023 was modest or flat for most companies, with only one in four increasing their spending by more than a quarter. That is set to change in 2024, with nine in ten respondents expecting to increase AI spending on data readiness (including platform modernization, cloud migration, and data quality) and in adjacent areas like strategy, cultural change, and business models. Four in ten expect to increase spending by 10 to 24%, and one-third expect to increase spending by 25 to 49%.


Data liquidity is one of the most important attributes for AI deployment. The ability to seamlessly access, combine, and analyze data from various sources enables firms to extract relevant information and apply it effectively to specific business scenarios. It also eliminates the need to sift through vast data repositories, as the data is already curated and tailored to the task at hand.

Data quality is a major limitation for AI deployment.

Half of respondents cite data quality as the most limiting data issue in deployment. This is especially true for larger firms with more data and substantial investments in legacy IT infrastructure. Companies with revenues of over US \$10 billion are the most likely to cite both data quality and data infrastructure as limiters, suggesting that organizations presiding over larger data repositories find the problem substantially harder.

Governance, security, and privacy are the biggest brake on the speed of AI deployment, cited by 45% of respondents.

Companies are not rushing into AI. Nearly all organizations (98%) say they are willing to forgo being the first to use AI if that ensures they deliver it safely and securely. Governance, security, and privacy are the biggest brake on the speed of AI deployment, cited by 45% of respondents (and a full 65% of respondents from the largest companies).



02 Partnering for success



Few companies are flying solo in the AI age. The cost and complexity of creating large language models (LLMs) and generative models from scratch is prohibitive and an abundance of platforms and tools are already on the market. “For most organizations, building their own large language models is very expensive and the value is time-limited,” says Kevin Collins, founder and CEO of Charli AI, an AI solutions provider. “If you don’t have the necessary expertise and resources to make a significant investment, it’s better to fine-tune and optimize off-the-shelf models.”


AI is also increasingly being integrated into existing software platforms, including those from giant technology providers such as Microsoft and Adobe. “We are infusing AI into everything we do within our organization and similarly taking that same mindset to our clients,” says SP Singh of Infosys. Companies, of course, can also pay to use proprietary models from providers such as OpenAI or fine-tune those models as needed. They can also build their own generative AI tools using open-source models.

Figure 1: Companies seek AI use cases tailored to industry needs


Is your organization developing the following types of AI use cases?
(Organizations that are developing AI use cases.)

- Use cases that are common across industries (e.g., generative AI for text generation or chatbots)
- Use cases that are specific to our industry (e.g., AI for drug discovery)
- Use cases that are unique to our business


Overall




Companies with annual revenue of \$1 billion to \$10 billion



Companies with annual revenue of \$500 million to \$1 billion



Companies with annual revenue more than \$10 billion



Matt McLarty, CTO at the software company Boomi, concurs that building LLMs will not be the path forward for most companies. “Most organizations are not going to build large language models, which are expensive and require a massive amount of infrastructure,” he says. “Instead, they need to become experts in applying the new technologies in their own business context. The companies that don’t put the AI cart before the business problem horse are going to be better positioned.”

Our survey results suggest that businesses are taking that advice, seeking AI use cases that address their unique business problems. While many firms are deploying general-purpose uses of AI (generative AI-powered chatbots for customer service, for example), a larger share say they are developing industry-specific use cases, or even use cases unique to their particular business (see Figure 1).

Sectoral nuances – including differing data formats, regulatory requirements, and usages – support the case for industry-tailored approaches to AI. From AI tools designed to document patient-doctor conversations to those leveraged to optimize drywall finishing, offerings tailored to the needs of specific industries or businesses are most likely to offer game-changing results.^{5,6}

Selecting a vendor

Companies have plenty of AI vendors to choose from; if anything, the choice can be daunting for executives who must balance cost, security, and diversification. “It feels overwhelming for companies – there is a lot of noise and questions around who to trust,” says Amy Machado, senior research manager at IDC, a market intelligence firm.

For privacy reasons, many are wary of relying on publicly available AI tools that may retain and reuse the data users enter, with some companies restricting employee use of certain LLMs.^{7,8} “If you’re shuffling your data off to a general-purpose AI, you have to do your due diligence on how they’re treating your data and what they’re doing with it behind the scenes,” says Collins. More specialized LLM-based tools are more likely to protect organizations’ proprietary data, come with support, and be consistently updated to address vulnerabilities.

Ultimately, executives may find that stitching together a “multi-AI” environment best meets their business needs and captures the distinctive value propositions of different providers.

Finance-friendly AI

Financial services firms operate in a highly regulated environment with high standards for data privacy, compliance, and security. “In the financial services sector, trust and security are paramount,” says Kevin Collins at Charli AI. “Our customers are major investment banks, hedge funds, and big banks that are analyzing S&P 500 companies. Everything we do needs to be trusted and fact-checked.”

In financial services, valuable information is often locked away in contracts, legal documents, PDF files, spreadsheets, and PowerPoint presentations. Extracting and integrating this data for analysis can be a daunting task. Specialized services like Charli AI have opened new possibilities for unlocking this trapped data: the company’s tools mine and curate a range of financial data, freeing up time for analysts to perform higher-value tasks.

“New AI has the ability to understand ‘locked-up’ content far better than ever before, so we’re unlocking a goldmine,” says Collins. “Our customers refer to us as the automated analysts for Wall Street – we can do in 30 minutes what it may take an analyst 80 hours of poring over a 500-page SEC filing to do.”

While companies are aware of the potential value of their data, many struggle to leverage it effectively. “They’re very data literate, but they also understand how challenging it is to get that data into a format that’s suitable for analytics and AI,” says Collins.

Our survey found that financial services respondents were least likely to be currently developing AI for industry-specific use cases, but companies like Charli are aiming to change this.

03 Counting the cost



From the electricity requirements for building and running AI models to the software, services, and salaries needed to manage solutions at scale, cost is a major consideration for AI. Google's US \$8 billion of spending in the third quarter of 2023 alone was overwhelmingly driven by AI.⁹ Training OpenAI's GPT-4 required US \$78 million of compute, while Google's Gemini Ultra had a compute cost of US \$191 million.¹⁰

Hardware is a major source of price uncertainty, including demand for graphics processing units (GPUs). "AI is very expensive, not just for the development and training of models, but also for the

operations and hardware costs, particularly GPUs," says Collins. Newer GPUs are more expensive, harder to obtain, and have a limited lifespan requiring a recurring capital cost. He adds, "AI is not cheap. A lot of our customers are shocked when they realize how much infrastructure they have to have to get the performance that they need."


While tech companies are largely shouldering hardware and computation spending in their race for market supremacy, firms inevitably face their own costs in areas like data management and monitoring, which are rarely built into lower-cost general-purpose LLMs. They must also invest to ensure compliance with

Figure 2: Changes in AI readiness spending from 2022 to 2023

How did your organization's investment in AI readiness change in the last year (from 2022 to 2023)?


Data-related investments

(platform modernization, cloud migration, pipelines, quality, etc.)



Investments in other areas

(developing AI strategy, cultural change, building models, etc.)



- Increased by more than 100%
- Increased by 50%–100%
- Increased by 25%–49%
- Increased by 10%–24%
- Stayed nearly the same
- Decreased



Source: MIT Technology Review Insights survey, 2024

Figure 3: Large companies boosted AI spending more aggressively

Companies with revenue greater than \$10 billion: How did your organization's investment in AI readiness change in the last year (from 2022 to 2023)?

Data-related investments
(platform modernization, cloud migration, pipelines, quality, etc.)

Investments in other areas
(developing AI strategy, cultural change, building models, etc.)



- Increased by more than 100%
- Increased by 50%–100%
- Increased by 25%–49%
- Increased by 10%–24%
- Stayed nearly the same
- Decreased

Source: MIT Technology Review Insights survey, 2024

firm-specific regulations, fund maintenance to ensure the systems they build remain up-to-date and relevant to the use case, and pay for the increased energy used by AI.¹¹ Talent costs include either hiring AI-skilled workers or upskilling the existing workforce.

Spending expectations

Polls show that larger firms, by revenue, are more likely to have AI in production, suggesting that financial constraints may be slowing smaller and medium-sized companies.¹² Analysis from CCS Insights suggests cost will contribute to a slowdown in generative AI adoption throughout 2024.¹³ As the hype cools, decision-makers may become more wary, especially as the productivity gains of generative AI may not be apparent for years.¹⁴

Such findings are corroborated by our survey, which found a link between company size and spending.

Budgetary constraints were the biggest obstacle to speed of AI deployment for medium-sized firms (annual revenue of US \$500 million to \$1 billion), with 47% citing this issue compared to a survey average of 22%. This indicates a “squeezed middle” of companies that need to adopt AI to stay competitive in their markets but for whom the costs are prohibitive.

To date, most firms have kept AI readiness spending flat, with larger companies more likely to have ramped up (see Figures 2 and 3).

However, nearly all companies expect to boost spending in the coming year, with 9 in 10 expecting to increase investment by at least 10% and one third anticipating spending up to 49% more both in data and adjacent areas like strategy and cultural change (see Figure 4).

“AI is not cheap. A lot of our customers are shocked when they realize how much infrastructure they have to have to get the performance that they need.”


Kevin Collins, Founder and Chief Executive Officer, Charli AI

Figure 4: AI readiness spending set to rise significantly into 2024


How will your organization's investment in AI readiness change in the next year (from 2023 to 2024)?

Data-related investments

(platform modernization, cloud migration, pipelines, quality, etc.)

**Investments in other areas**

(developing AI strategy, cultural change, building models, etc.)



- Will increase by more than 100%
- Will increase by 50%–100%
- Will increase by 25%–49%
- Will increase by 10%–24%
- Will stay nearly the same
- Will decrease

Source: MIT Technology Review Insights survey, 2024

As companies dial in their AI investments, AI cost dynamics will of course also evolve. Innovation will doubtless bring down costs through efficiency gains. The cost curve for computation will also fall as startups and large companies address the shortcomings of existing hardware and models.

Measuring return on investment

To weigh the financial outlays of AI, organizations need to develop robust return on investment (ROI) methodologies capturing not just the efficiency gains of automating “business as usual” tasks, but also the new value AI can create.

Companies are already quantifying AI’s cost savings. Motorola has developed a framework that tracks hours taken to complete a task with and without generative AI.¹⁵ Swedish fintech company Klarna has already calculated generative AI’s output as equivalent to 700 customer service agents.¹⁶ But “the mindset is shifting to using AI as an enabler for revenue growth, not just cost savings,” says Machado.

As employees become more productive, new opportunities emerge that may be hard to quantify. For example, AI allows data scientists to explore new ideas and expand into new products. The automation of routine tasks by copilots can alleviate workloads and

“The mindset is shifting to using AI as an enabler for revenue growth, not just cost savings.”

Amy Machado, Senior Research Manager, IDC

reduce burnout, allowing employees to dedicate more time to more creative and innovative projects.

As enterprise use of AI spreads, employee expectations will shift and AI-enhanced workplace experiences will become necessary to attract talent. AI professionals already have different expectations than traditional job seekers, ranking interesting job content and work on cutting-edge projects higher among their job priorities.¹⁷

The consultancy PwC compares “hard ROI,” such as time saved and productivity boosts, with “soft ROI,” including improved employee experience and ability to respond flexibly to new opportunities.¹⁸ Accurately assessing AI’s ROI requires metrics beyond traditional financial measures. For instance, Collins points out that tools like Charli AI’s unlock huge value by helping companies make sense of their data more efficiently.


04

Building a data core

Data quality, infrastructure, and governance are all essential to deliver AI workloads efficiently and at scale. While AI technology can be transformative, without the data foundations in place, organizations will struggle to both collect the right data and derive insight from it. “People still have the perception that AI today is magic or super intelligent, and that’s far from the truth,” says Collins. “AI is a science and a tool. You still have to do all of the hard work around data governance and figuring out the right application of these tools.”

Organizations that overlook these supporting dimensions may struggle to deploy AI successfully. Data quality problems, cited by half of respondents as their most limiting data issue, can severely hinder AI performance. “Going forward, the conversation will be around getting data AI-ready and in the right format to actually utilize it effectively,” says Machado. Data infrastructure or pipelines, data integration tools, and cloud migration were also mentioned by many respondents as major barriers to deployment, indicating the importance of a robust IT architecture.

These issues are particularly prominent for companies with revenues of over US \$10 billion, which are the most likely to cite both data quality (52%) and data infrastructure (55%) as obstacles, compared with overall survey averages of 49% and 44% (see Figure 5). Organizations presiding over larger data repositories and legacy IT infrastructure may be finding greater complexity and costs in transitioning to an AI-ready architecture.



Data management: Tips and tactics


- Track and manage the lineage of your organization's data to maintain data quality and integrity as it moves from one analytical model or application to the next.
- Optimize communication across the organization on AI, and on data aggregation and requirements, ensuring cross-functional teams collaborate on business cases.
- Instead of focusing on data centralization, adopt a strategy that prioritizes data contextualization, such as extracting metadata and storing it in an appropriate metadata language.
- Identify your organization's core capabilities and competencies, ensuring you are able to adapt to rapid changes and are not overly reliant on legacy infrastructure.

“AI is a science and a tool. You still have to do all of the hard work around data governance and figuring out the right application of these tools.”

Kevin Collins, Founder and Chief Executive Officer, Charli AI

Figure 5: Data bottlenecks in AI deployment

Which aspects of your organization's data are most limiting the speed to deploy AI?



Source: MIT Technology Review Insights survey, 2024

Reckoning with legacy infrastructure

Transitioning away from legacy IT infrastructure can be daunting for organizations of all sizes. “Legacy systems and sunk costs are definitely an inhibitor for all kinds of data management initiatives,” says Machado. “Many organizations don’t want to touch certain legacy systems because ‘it’s kind of working’ and only one person knows how to maintain that older technology.” These legacy systems are often based on fragmented, outdated architectures and programming languages that are hard to integrate with modern AI solutions.

To reap the benefits of AI, organizations must develop an IT architecture capable of accommodating both structured and unstructured data across the entire lifecycle, from the source to processing, analytics, and storage. Unstructured data can yield considerable insight for AI systems, but its raw information can be difficult to assimilate into existing systems. That can be problematic since, according to Machado, 90% of the data in an enterprise content management system is unstructured.

A centralized repository solution in the form of a data lake can be useful, but it must be well organized and structured to prevent it becoming filled with irrelevant and unusable data. “It’s not enough to just have a big data lake, because that’s going to become a data swamp; you have to have good governance over the data lineage and how it flows from one model to the next,” says Collins.

Nearly half (46%) of survey respondents also cite manual, non-digitized, or non-automated processes for data handling and management as a limiting factor for AI deployment. Without strategies to fully digitize processes across the business, organizations will leave valuable data sources untapped and limit AI’s automation potential.

Data lineage and liquidity

Deciding which data to leverage requires it to be recombined and contextualized. Implementing data lineage tracking and quality-assurance measures can support this contextualization process, preserving the integrity of data as it undergoes processing and analysis.

According to McLarty, data liquidity, or the ability to seamlessly access, combine, and analyze data from various sources, is one of an organization’s most important AI assets. Data liquidity empowers organizations to understand the relationships, dependencies, and nuances within their data, enabling them to extract relevant information and apply it effectively to specific business scenarios. “Organizations with high data liquidity – the ability to get the right data at the right time and place – will be most successful with AI,” says McLarty.

Data liquidity helps ensure data is readily accessible, usable, and shareable. Nurturing liquid data assets requires a flexible, harmonized and scalable IT infrastructure that can seamlessly handle increasing volumes and varieties of data without disrupting existing operations.

Metadata

Extracting metadata – or data about data – helps to contextualize data for use in AI models. “Organizations want to understand their data upfront to see what is valuable for real business purposes, and metadata holds everything together to make this possible,” says Machado.


Within a content management system, data sources from different systems can be connected by a common metadata language, and it is important, Machado says, to select a language suitable for the organization’s use case. For sharing information on biological specimens, for instance, the Darwin Core language might be considered appropriate. VRA Core is used for works of visual culture.

By curating metadata, enterprises can harness the power of their previously untapped unstructured data reserves. This approach allows them to optimize their data for AI solutions that will drive business value.

“Organizations with high data liquidity – the ability to get the right data at the right time and place – will be most successful with AI.”

Matt McLarty, Chief Technology Officer, Boomi

05 Acceleration versus caution



A significant barrier to the rapid adoption of AI is the need to ensure robust governance, security, and privacy measures are in place, especially in the case of new risks posed by generative AI. These issues serve as the biggest brake on the speed of AI deployment for 45% of companies – and for 65% of companies with over \$10 billion in revenue. That difference may be due to the increased risk profile of larger organizations, including their larger “attack surfaces,” making them more vulnerable to potential breaches and regulatory scrutiny.

Hallucinations, errors, and bias

Generative AI comes with a range of new risks, including producing inaccurate information or “hallucinations”; legal risks such as plagiarism, copyright infringement, and liability for errors; challenges around privacy and data ownership; and systemic bias. These risks can harm customers and expose companies to legal liability. Air Canada was recently held liable for damages caused by misleading information provided by an AI chatbot.¹⁹

Inaccurate information and hallucinations can contribute to heightened bias in AI models. In financial services, for instance, AI-powered credit risk assessments could help mitigate systemic discrimination in credit applications – but only if the data used to train these AI systems is itself reliable, accurate, and free of distortions and biases. An AI-powered solution built on discriminatory data and algorithms, however, will amplify existing biases,

resulting in even less fair outcomes. Similar biases can crop up in applications including health care, policing, and advertising, where algorithms can accentuate gender biases or overlook underrepresented minority groups.

Cyber risk

The use of generative AI tools also comes with new opportunities for cyberattacks. In a prompt injection attack, bad actors craft prompts to elicit unintended responses from an LLM, to manipulate its behavior, achieve unauthorized access, or bypass security measures. In February 2023, a student at Stanford University discovered sensitive details about Bing Chat using this technique, instructing it to “ignore previous instructions” and reveal hidden information.²⁰


In data poisoning, the data used to train an AI model is intentionally targeted and injected with malicious information. Leakages of data from databases and pipelines that feed LLMs are another source of concern.

Data privacy and protection

Data privacy and protection demand significant investment in the AI age. This is of particular concern for high-risk sectors and for those to whom security and regulatory compliance are paramount. In our survey, respondents from financial services were 13 points more likely to cite governance, security, and privacy as factors inhibiting AI adoption (58% of respondents, compared to the survey average of 45%).

Figure 6: Executives willing to delay AI use to assure safety and security

Respondents who agree with the following statement: “We are willing to forego being first to use AI in our sector if that ensures that we deliver it safely and securely.”



Source: MIT Technology Review Insights survey, 2024

Our survey finds that nearly all companies across sectors are taking safety and security seriously, with a remarkable 98% of respondents saying they are willing to forego first-mover advantage if it ensures AI is delivered safely and securely (see Figure 6). This cautious approach highlights the widespread recognition that AI risks can be potentially devastating to brand reputation or prompt a stringent regulatory response. Companies are taking the long view, striking a balance between progress and security.

In many cases, firms are leveraging AI itself to strengthen their cybersecurity and resilience. Companies now must account for threats at a “machine scale, rather than a human scale,” and organizations see the power of AI solutions for this work.²¹ AI tools can participate in vulnerability scanning, risk assessments, threat detection, and proactive threat hunting.

Despite the progress made, every new AI solution can also create new threats. Human oversight will remain an important part of the process, not least because the human element can help improve these tools. “Given trust and governance concerns, we’re still in a position where we need to have lots of human guardrails in the picture,” says McLarty.

Rising regulatory tide

The regulatory response to AI thus far has focused on establishing safeguards and governance frameworks to supervise AI activity, control risk, and provide a level playing field. An analysis of 127 countries shows that the number of bills mentioning “artificial intelligence” passed into law worldwide grew from just 1 in 2016 to 37 in 2022.²² While this reflects a recognition of the transformative impact of AI, there is still a need to accelerate the implementation of comprehensive policies and safeguards on ethics, privacy and security.

The European Union’s AI Act is one of the most comprehensive initiatives currently being enacted. The Act establishes a risk-based approach to categorizing AI systems and imposes stricter requirements on systems used in higher-risk areas such as infrastructure, employment, and law enforcement. AI used in these areas will face rigorous obligations related to data quality, transparency, human oversight, accuracy, and other considerations before being allowed on the market. Companies implementing high-risk systems will need to complete a conformity assessment before being placed on the market, with low-risk systems such as chatbots subject to voluntary codes of conduct and transparency requirements.²³

Nearly all companies across sectors are taking safety and security seriously, with a remarkable 98% of respondents saying they are willing to forego first-mover advantage if it ensures AI is delivered safely and securely.

Certain uses of AI, such as social scoring and mass surveillance, meanwhile, will be banned outright.²⁴

Similar initiatives are underway in the US, where President Joe Biden issued an executive order setting out new standards for AI security and safety in October 2023. The order imposes a degree of human oversight through measures such as an advanced cybersecurity program aimed at developing AI tools to find and fix vulnerabilities in critical software. It also establishes an AI Safety and Security Board, and it requires that developers of the most powerful AI systems share safety test results with the government.²⁵

Compliance challenges

The flurry of activity around AI laws and policies will be an ongoing challenge for enterprise. “It can take an organization a couple of years to get up to speed on how to comply with new regulations, given the amount of reading required to understand all the policies and guidelines,” says Collins.

A key challenge will be establishing workable systems for assessing compliance and performance. AI assurance processes such as algorithm impact assessments and bias audits help organizations monitor and evaluate the capabilities of their AI systems, playing a crucial role in building confidence and trust.²⁶ Gaps in regulation make devising frameworks for AI assurance challenging, but standards like those set by the International Organization for Standardization (ISO) offer guidance.²⁷

The current scope of AI regulation recognizes that risk can never be avoided entirely. Instead, the regulation is centered on disincentivizing irresponsible and risky AI deployment. Going forward, savvy organizations may prioritize AI explainability – the ability to understand the decision-making process and outputs of an AI model – both for compliance reasons and to build customer trust.



06

Conclusion



Over the last year, organizations have come to understand the power and potential of AI. This year, those organizations are seeking to shift from small pilots to enterprise-wide deployment of AI at scale.

This report aims to assist with the next steps of that journey, offering the following principles as guidance.

This is a critical year for building AI foundations. As organizations strive to meet ambitious short-term AI goals, they are discovering that having the right data foundations is critical. Organizations are responding by doubling down on investments in data quality, data liquidity, and IT infrastructure. “There’s no question in my mind,” says McLarty, “that companies that have good hygiene and rigor around their data are going to be way better positioned for the AI landscape.”

The highest-leverage AI use cases drive targeted, business-specific results. Although general-purpose generative AI use cases are increasingly simple to implement, they are equally available to competitors and customers. The most valuable use cases are those that create unique competitive advantage for the business. “You can deploy AI capabilities, but if they don’t provide true business value, you’re just wasting money to say

you did something cool, with no real reason. That undermines trust in future projects,” warns Machado.

Financial considerations and partnerships are key. This year, the costs of AI – from GPUs to skilled talent to energy consumption – must be accounted, and a realistic approach to measuring AI’s ROI must be developed. Because few organizations will go it alone when it comes to AI, making decisions about the right partners, vendors, and tools will be essential to most companies’ AI futures.

Speed of AI adoption is being moderated by a realistic assessment of its risks. Organizations are rightly cautious about the risks posed by injudicious use of AI, and essentially all agree that caution is preferable to first-mover advantage when scaling AI. An emerging regulatory environment and a better understanding of how the risks of AI can be mitigated should serve to speed adoption.

These principles, it should be noted, are not deeply technological. “The biggest factor holding back AI implementation,” says McLarty, “is people not knowing where to start. People need to understand that you don’t actually have to be an expert on how to create generative AI in order to get value from generative AI.”

“The biggest factor holding back AI implementation is people not knowing where to start. People need to understand that you don’t actually have to be an expert on how to create generative AI in order to get value from generative AI.”

Matt McLarty, Chief Technology Officer, Boomi

Endnotes

1. "Sizing the prize: What's the real value of AI for your business and how can you capitalise?" PwC, 2017, <https://preview.thenewsmarket.com/Previews/PWC/DocumentAssets/476830.pdf>.
2. "40% of time spent on mundane chores could be automated within 10 years – AI experts," University of Oxford, February 22, 2023, <https://www.ox.ac.uk/news/2023-02-22-40-time-spent-mundane-chores-could-be-automated-within-10-years-ai-experts-0>.
3. "AI investment forecast to approach \$200 billion globally by 2025," Goldman Sachs, August 1, 2023, <https://www.goldmansachs.com/intelligence/pages/ai-investment-forecast-to-approach-200-billion-globally-by-2025.html>.
4. Richard Vanderford, "Generative AI Isn't Ubiquitous in the Business World – at Least Not Yet," The Wall Street Journal, April 7, 2024, https://www.wsj.com/articles/generative-ai-isnt-ubiquitous-in-the-business-world-at-least-not-yet-07c07f01?mod=tech_lead_pos3.
5. "AWS Announces AWS HealthScribe," AWS, July 26, 2023, <https://press.aboutamazon.com/2023/7/aws-announces-aws-healthscribe-a-new-generative-ai-powered-service-that-automatically-creates-clinical-documentation>.
6. "Canvas + USG Partnership Aims to Transform Drywall Industry," Canvas, May 11, 2023, <https://www.canvas.build/canvas-usg-partnership-aims-to-transform-drywall-industry/>.
7. Aaron Tilley and Miles Kruppa, "Apple Restricts Employee Use of ChatGPT, Joining Other Companies Wary of Leaks," The Wall Street Journal, May 18, 2023, <https://www.wsj.com/articles/apple-restricts-use-of-chatgpt-joining-other-companies-wary-of-leaks-d44d7d34>.
8. Siladitya Ray, "Samsung Bans ChatGPT Among Employees After Sensitive Code Leak," Forbes, May 2, 2023, <https://www.forbes.com/sites/siladityaray/2023/05/02/samsung-bans-chatgpt-and-other-chatbots-for-employees-after-sensitive-code-leak/?sh=34b316360781>.
9. Jennifer Elias, "Wall Street wants to know how Google's going to profit from AI," CNBC, October 24, 2023, <https://www.cnbc.com/2023/10/24/wall-street-wants-to-know-how-googles-going-to-profit-from-ai.html>.
10. "Artificial Intelligence Index Report 2024," Stanford University, 2024, https://aiindex.stanford.edu/wp-content/uploads/2024/04/HAI_AI-Index-Report-2024.pdf.
11. Ryan Peeler, "The Hidden Costs of Implementing AI In Enterprise," Forbes, August 31, 2023, <https://www.forbes.com/sites/forbestechcouncil/2023/08/31/the-hidden-costs-of-implementing-ai-in-enterprise/?sh=932cbac4d5c1>.
12. "2023 Global Trends in AI Report," S&P Global, 2023, <https://www.weka.io/resources/analyst-report/2023-global-trends-in-ai/>.
13. "CCS Insight's Predictions for 2024 and Beyond Revealed," CCS Insight, October 10, 2023, <https://www.ccsinsight.com/company-news/ccs-insights-predictions-for-2024-and-beyond-revealed/>.
14. "Gartner Places Generative AI on the Peak of Inflated Expectations on the 2023 Hype Cycle for Emerging Technologies," Gartner, August 16, 2023, <https://www.gartner.com/en/newsroom/press-releases/2023-08-16-gartner-places-generative-ai-on-the-peak-of-inflated-expectations-on-the-2023-hype-cycle-for-emerging-technologies>.
15. Kristin Broughton and Mark Maurer, "CFOs Tackle Thorny Calculus on Gen AI: What's the Return on Investment?" March 24, 2024, https://www.wsj.com/articles/cfos-tackle-thorny-calculus-on-gen-ai-whats-the-return-on-investment-24ebf435?mod=series_proaiimplementation.
16. Richard Vanderford, "Generative AI Isn't Ubiquitous in the Business World – at Least Not Yet," The Wall Street Journal, April 7, 2024, https://www.wsj.com/articles/generative-ai-isnt-ubiquitous-in-the-business-world-at-least-not-yet-07c07f01?mod=tech_lead_pos3.
17. "How to Attract, Develop, and Retain AI Talent," BCG, 2023, <https://www.bcg.com/publications/2023/how-to-attract-develop-retain-ai-talent>.
18. Anand Rao, "Solving AI's ROI problem. It's not that easy," PwC, July 20, 2021, <https://www.pwc.com/us/en/tech-effect/ai-analytics/artificial-intelligence-roi.html>.
19. Maria Yagoda, "Airline held liable for its chatbot giving passenger bad advice – what this means for travellers," BBC, February 23, 2024, <https://www.bbc.com/travel/article/20240222-air-canada-chatbot-misinformation-what-travellers-should-know>.
20. Jacob Fox, "Prompt Injection Attacks: A New Frontier in Cybersecurity," Cobalt, May 31, 2023, <https://www.cobalt.io/blog/prompt-injection-attacks>.
21. "Cisco AI Readiness Index," Cisco, 2023, https://www.cisco.com/c/dam/m/en_us/solutions/ai/readiness-index/documents/cisco-global-ai-readiness-index.pdf.
22. "Artificial Intelligence Index Report 2023," Stanford University, 2023, https://aiindex.stanford.edu/wp-content/uploads/2023/04/HAI_AI-Index-Report_2023.pdf.
23. "The AI Act of the European Union," TaylorWessing, <https://www.taylorwessing.com/en/insights-and-events/insights/artificial-intelligence-act>.
24. "What Is The EU Artificial Intelligence Act?" European Movement Ireland, September 18, 2023, <https://www.europeanmovement.ie/what-is-the-eu-ai-act/>.
25. "FACT SHEET: President Biden Issues Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence," The White House, October 30, 2023, <https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence>.
26. "Introduction to AI assurance," UK Government, February 12, 2024, <https://www.gov.uk/government/publications/introduction-to-ai-assurance/introduction-to-ai-assurance>.
27. "AI management systems: What businesses need to know," International Organization for Standardization, <https://www.iso.org/artificial-intelligence/ai-management-systems>.

About MIT Technology Review Insights

MIT Technology Review Insights is the custom publishing division of *MIT Technology Review*, the world's longest-running technology magazine, backed by the world's foremost technology institution – producing live events and research on the leading technology and business challenges of the day. Insights conducts qualitative and quantitative research and analysis in the US and abroad and publishes a wide variety of content, including articles, reports, infographics, videos, and podcasts. And through its growing MIT Technology Review **Global Insights Panel**, Insights has unparalleled access to senior-level executives, innovators, and entrepreneurs worldwide for surveys and in-depth interviews.

About Boomi

Boomi, the intelligent integration and automation leader, helps organizations around the world automate and streamline critical processes to achieve business outcomes faster. Harnessing advanced AI capabilities, the Boomi Enterprise Platform seamlessly connects systems and manages data flows with API management, integration, data management, and AI orchestration in one comprehensive solution. With a customer base exceeding 20,000 companies globally and a rapidly expanding network of 800+ partners, Boomi is revolutionizing the way enterprises of all sizes achieve business agility and operational excellence. Discover more at [boomi.com](https://www.boomi.com).




Illustrations

Illustrations assembled by Tim Huxford with elements from Shutterstock and Adobe Stock.

While every effort has been taken to verify the accuracy of this information, MIT Technology Review Insights cannot accept any responsibility or liability for reliance by any person in this report or any of the information, opinions, or conclusions set out in this report.

© Copyright MIT Technology Review Insights, 2024. All rights reserved.



MIT Technology Review Insights

www.technologyreview.com

insights@technologyreview.com