



From compliance to confidence

Embracing a new mindset to advance responsible AI maturity

>
accenture

Contents

Page 03

Executive summary:
rethinking responsible AI

Page 12

Becoming reinvention
ready: the milestones of
responsible AI maturity

Page 05

Risky business:
the three challenges of
the current risk landscape

Page 15

Responsibility reality check:
how ready are companies
for responsible AI?

Page 10

From compliance to
value: companies are
acknowledging the impact
of responsible AI

Page 22

Ready, set, grow:
five priorities for
responsible AI



Executive summary: rethinking responsible AI

As generative AI becomes pervasive in business and society, so too do the risks that come with using it. Consider the chatbot that gave incorrect advice to a customer, creating liability for the company that launched the chatbot. Or the employee who accidentally exposed proprietary company data after using ChatGPT. Or the algorithm that wrongly flagged thousands of individuals for fraud.

In this new environment, the responsible development and use of data and AI becomes a critical enabler for organizations to both minimize the risks that come with using these technologies and to unlock its many benefits—from helping individuals perform tasks better and faster to allowing companies to reinvent themselves to create more value and gain a competitive edge through innovation.

This is, of course, easier said than done, for the challenges of scaling responsible AI across an organization are daunting. AI-related risks continue to accumulate as generative AI creates and accelerates both data and AI risks. New AI-focused laws and regulations are germinating everywhere. And AI value chains¹ are growing in complexity, especially as more companies become both developers and buyers of AI models.

Responsible AI 101

What is responsible AI?

Responsible AI implies taking intentional actions to design, deploy and use AI to create value and build trust by protecting against the potential risks of AI.

What is mature responsible AI?

Having fully operationalized responsible AI efforts as a platform to take a more systemic, future-orientated approach that unlocks the true value of AI. The North Star for responsible AI maturity is to become a pioneer. No companies are yet at this stage.

How do you build a responsible AI program?

Our research and work advising clients has shown that all companies can benefit from focusing on these five priorities to improve their maturity and begin to reap the benefits of AI.

01

Establish AI governance and principles

02

Conduct AI risk assessments

03

Systemic enablement for responsible AI testing

04

Ongoing monitoring and compliance

05

Workforce impact, sustainability, privacy, security

Our survey respondents estimate that when a company becomes a pioneer in responsible AI, its AI-related revenue will increase by

18% on average.

Today's era of generative AI creates new requirements when it comes to responsible AI. To be a leader in responsible AI, companies must pursue an anticipatory mindset, commit to continuous improvement and extend their focus beyond their organization to their entire value chain and wider AI ecosystem. We call this new level of maturity "becoming a responsible AI pioneer", and no companies are yet at this stage.

To better understand companies' attitudes toward AI-related risks, as well as their approach to responsible AI, we collaborated with Stanford University to

survey C-suite executives across 1,000 companies, spanning 19 industries and 22 countries.²

We assessed companies' maturity in responsible AI by developing a four-stage framework—the higher the stage, the greater the progress. We then applied that framework to analyze the organizational and operational maturity of the 1,000 companies we surveyed.

To their credit, most business leaders recognize the importance of responsible AI in unlocking business value.

Responsible data is an integral part of the responsible AI journey

It's important to note that when we talk about responsible AI, the responsible use of data is a critical part of the equation. The importance of data in the age of generative AI cannot be underestimated. Most large organizations are still grappling with longstanding issues around data quality, availability and governance. In a separate global survey by Accenture of 2,000 business leaders in 2024, 48% said their organizations lacked enough high-quality data to operationalize their generative AI initiatives.³ Companies with high data readiness have the right data, with the right quality and quantity. They have a scaled data management and governance practice in place which allows use of data seamlessly across their processes, ensuring responsible adoption and enabling the economic value of data.

[Click here](#) to learn more about key actions companies can take to improve their data readiness in the era of generative AI.

They estimate that when a company becomes a pioneer in responsible AI, its AI-related revenue will increase by 18% on average. However, our research found that a great majority of companies we surveyed are not as prepared for responsible AI as they would like.

Years of client experience show us that companies can take action to systemically operationalize responsible AI. In this report, we will explore the actions that will provide the organizational infrastructure to become a responsible AI pioneer and unlock the true value of AI.

Risky business: the three challenges of the current risk landscape

AI-related risks are evolving quickly, especially as generative AI spreads far and wide. In the last 12 months alone, we have seen a new wave of risks emerge (such as those connected to hallucinations, intellectual property (IP), cybersecurity and environmental impact) while more established risks like data privacy, reliability and transparency take on renewed prominence. This creates an increasingly complex risk landscape.

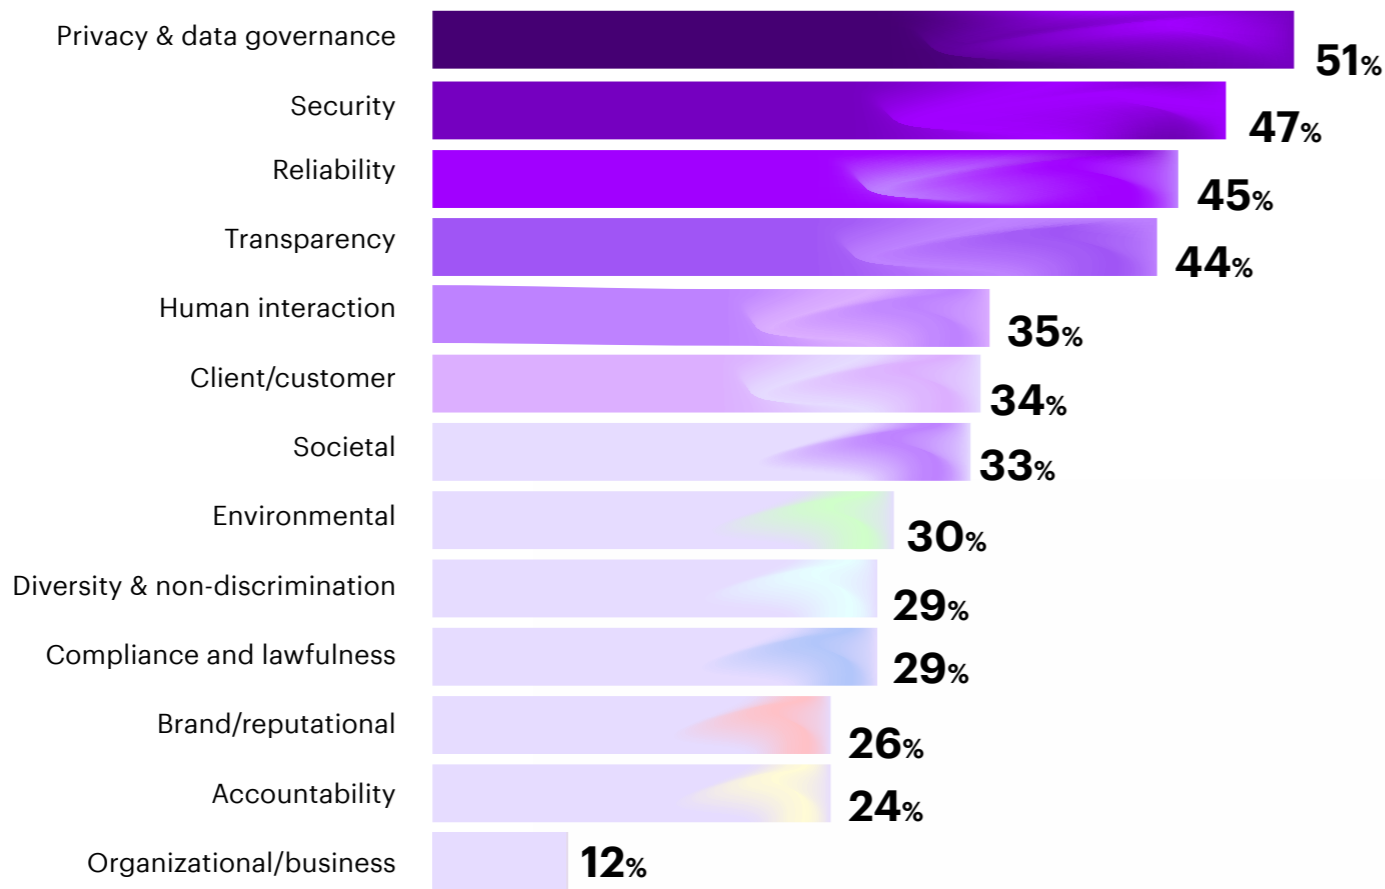
We found that the risk landscape will continue to expand and evolve across three main areas.

01 An increasing range and frequency of risk	02 A continuously evolving regulatory landscape	03 An expanding scope of risk management across the value chain
---	--	--

01 An increasing range and frequency of risk

The most-cited risk, as Figure 1 shows, involved concerns about privacy and data governance (51% of respondents cited this as a risk for their company). Security (cited by 47% of respondents) and reliability risks, such as output errors, hallucinations and model failure (a concern for 45% of respondents) were second and third, respectively.

Figure 1: Top AI-related risks for companies



Source: Accenture Stanford Executive Survey, N=1,000. See "About the research" for details.

Generative AI’s impact on the evolving AI risk landscape is also visible in many of the risks that worry companies most. These touch on everything from transparency and the challenge of “black box” models (cited by 44% of respondents), to the environment and the energy demands of data centers (a concern for 30% of respondents), to accountability and concerns about infringing on copyright and other intellectual property (cited by 24% of respondents). To learn more about the risk dimensions we assessed, [click here](#) to read the paper we produced in collaboration with Stanford University.

As the number and type of AI risks have increased sharply in recent years, AI-driven incidents (bias, deepfakes, hallucinations, privacy breaches etc.) have become much more common, increasing 32.3% in 2023, according to the AI Incident Database, a website that monitors such occurrences⁴. This trend should raise red flags for companies. The complexity of AI risk continues to expand in terms of range and frequency, accelerated by generative AI advances. As a result, responsible AI mitigation strategies must also evolve and change, so that companies can continue to adopt AI at speed. In this dynamic environment, companies can no longer react to risks, they must learn to anticipate them.



02 A continuously evolving regulatory landscape

It's also true that companies that fall behind on responsible AI will expose themselves to a growing risk of non-compliance, as more governments begin to regulate artificial intelligence.

The European Union's AI Act started the engines, with other AI regulation currently under consideration across more than 37 countries. As governments inevitably choose to regulate AI in different ways, compliance will grow increasingly difficult for multinational firms. This reality is already impacting most of the companies we surveyed, with **77% either facing AI regulation already or expecting to be subject to it over the next five years.**

As governments tend to regulate AI in different ways, complexity and confusion will likely grow for multinational companies.

What's more, it's not simply AI regulation organizations have to contend with. The emergence of new generative AI risks is forcing governments to take action by passing and amending existing laws,

which only adds to the complexity.

In fact, nearly all **(90%) of the companies surveyed expect to be subject to AI-adjacent legal obligations such as cybersecurity, AI-related laws, data and consumer protection over the next five years.** The fact that such legislation is happening at national and sub-national levels is creating further compliance challenges for companies.

For example, new IP or copyright laws have been proposed or introduced in China, Singapore, Brazil and Saudi Arabia. In 2023, South Korea introduced an AI liability law, which precedes a similar law planned in the EU, while the state of California amended the California Privacy Right Act (CPRA), giving consumers the right to opt out of practices such as the selling and sharing of personal data. Deepfake laws have been proposed or introduced in France, UK, Australia and South Korea, with a number of countries investigating new AI cybersecurity laws.

77%

of companies we surveyed either already face AI regulation or expect to be subject to it over the next five years.

90%

of companies surveyed expect to be subject to AI-adjacent legal obligations such as cybersecurity, AI-related laws, data and consumer protection over the next five years.



Global AI regulatory initiatives (as of publication)

~Listed initiatives are non-exhaustive~

UK

- Artificial Intelligence Regulation White Paper 2023
- Generative Artificial Intelligence in Education 2023
- Frontier AI Taskforce 2023
- UK AI Safety Institute 2023
- UK AI Safety Summit and Bletchley Declaration 2023
- US-UK Partnership on Science of AI Safety 2024
- UK-Canada Partnership on Science of AI Safety 2024
- [Regulating AI: the ICO's Strategic Approach 2024](#)

Canada

- Bill C-27: digital charter implementation act to include artificial intelligence and data act (AIDA part C) 2022
- Artificial Intelligence and Data Act (AIDA) 2023
- \$2.4B Federal investment 2024
- [Canada AI Safety Institute 2024](#)

France

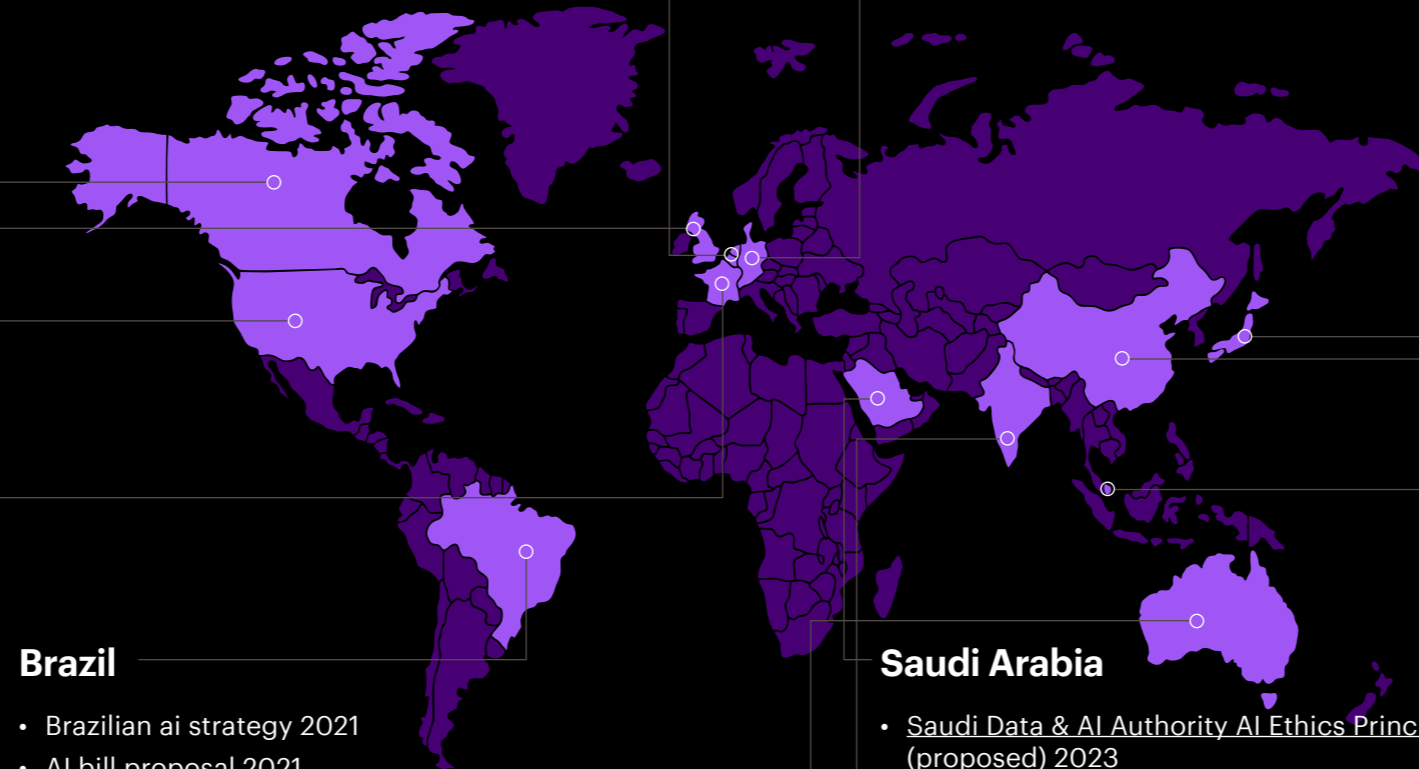
- [Conseil d'etat approach on AI governance 2022](#)
- [France AI Commission Report 2024](#)
- [France genAI plan 2024](#)

United States

- White House voluntary commitments from 15 leading tech companies on certain security and transparency requirements 2023
- Federal Election Commission - proposed regulations on the use of deceptive AI in campaign ads 2023
- National Telecommunications and Information Administration (NTIA) RFI on AI accountability 2023
- Executive Order on the Safe, Secure and Trustworthy Development and Use of Artificial Intelligence 2023
- NIST Draft GenAI Companion to RMF 2024
- [AI Safety Institute 2024](#)
- [National AI Research Resource \(NAIRR\) 2024](#)

EU

- Artificial intelligence liability directive 2022
- Data Governance Act 2022
- Digital Markets Act 2023
- EU AI Act 2024
- [Digital Services Act 2024](#)



Brazil

- Brazilian ai strategy 2021
- AI bill proposal 2021
- AI Bill No. 2238 2023 (proposed)
- [National Digital Government Strategy 2024](#)

Australia

- Safe and Responsible AI (interim response) (2024)
- National Framework for the Assurance of Artificial Intelligence in Government (2024)
- Policy for the responsible use of AI in government (2024)
- [Mandatory guardrails for AI in high-risk settings \(proposed\) \(2024\)](#)

Germany

- Fundamentals on security of AI systems 2022
- AI Action Plan 2023
- Testing and Certification Center for AI-based Robots 2024
- Guidelines for the use of AI in administrative and professional tasks 2024

Saudi Arabia

- [Saudi Data & AI Authority AI Ethics Principles \(proposed\) 2023](#)
- Saudi Data & AI Authority - Generative Artificial Intelligence Guidelines for Government January 10, 2024
- Saudi Data & AI Authority Generative Artificial Intelligence Guidelines for the Public January 11, 2024

India

- Operationalizing principles for responsible AI 2021
- Digital Personal Data Protection Act 2023
- Approval of AI tools before public release 2024
- India AI report 2023
- India AI Mission Framework 2024
- Recommendations on Encouraging Innovative Technologies Through Regulatory Sandbox. 2024

China

- Personal Information Protection Law 2022
- Internet Information Service Algorithm Recommendation Management Regulations 2022
- China's Deep Synthesis Provisions 2023
- [Draft Measures for the Management of Generative AI Services 2024](#)

Japan

- Social principles of human-centric AI 2019
- AI utilisation guidelines 2019
- AI governance guidelines 2022
- AI strategy 2022
- AI Strategy Council 2023
- Hiroshima Process 2023
- Japan AI Safety Institute 2024
- Legislation to regulate generative AI (proposed)
- international framework for regulating and using generative AI 2024

Singapore

- AI governance approach + implementation self-assessment guide 2020
- AI governance testing framework minimum viable product (mvp) 2021
- A.I. Verify 2022
- Singapore National AI Strategy 2.0 (NAIS 2.0) 2023
- AI safety labels (proposed) 2024
- AI funding initiative to power Singapore's economic growth (2024)
- Model AI Governance Framework for Generative AI (2024)
- Safety Guidelines for Model Developers and App Developers (2024)
- [MAS principles to promote FEAT in the use of AI](#)



03 An expanding scope of risk management across the value chain

It's not just the types of risks and regulations that are changing. As firms begin buying and developing AI models, they also need to prepare for the risks that come from doing both. A developer, for instance, may be especially worried about being sued when training a model in a way that violates IP laws. A buyer, on the other hand, may be more worried about whether a newly acquired AI model will perform as advertised. When a company becomes both a buyer and developer of models, the types of risks that it faces can get exponentially complex.

This double-sided risk represents another big shift in the AI landscape. Until recently, few companies both bought and developed AI models.

Today, however, almost one-third (28%) of the companies we surveyed take on the role of both buyer and developer (33% of firms just develop and 38% just buy).

Our research also suggests that many companies are insufficiently prepared for the evolving complexity of their AI value chains. For example, just 43% of surveyed firms that are acquiring AI models have robust procurement measures in place, such as regulatory checks and third-party audits. Organizations must work now to assess whether third-party AI products or services meet the organization's AI standards and are monitored to enable ongoing compliance and risk management.

As third-party risks continue to evolve, companies must think beyond their own responsible AI strategy. The reality is that companies must do their due diligence and make sure all legal and regulatory responsibilities are agreed on and met by all parties along the value chain. For high-risk AI use cases, it will not be sufficient to rationalize outcomes based on an "unintended consequence". Companies should expect to be held accountable by their customers and regulators in their oversight of high-risk use cases.

Only

43%

of companies acquiring AI models have procurement measures in place, such as regulatory checks and third-party audits.



From compliance to value: companies are acknowledging the impact of responsible AI

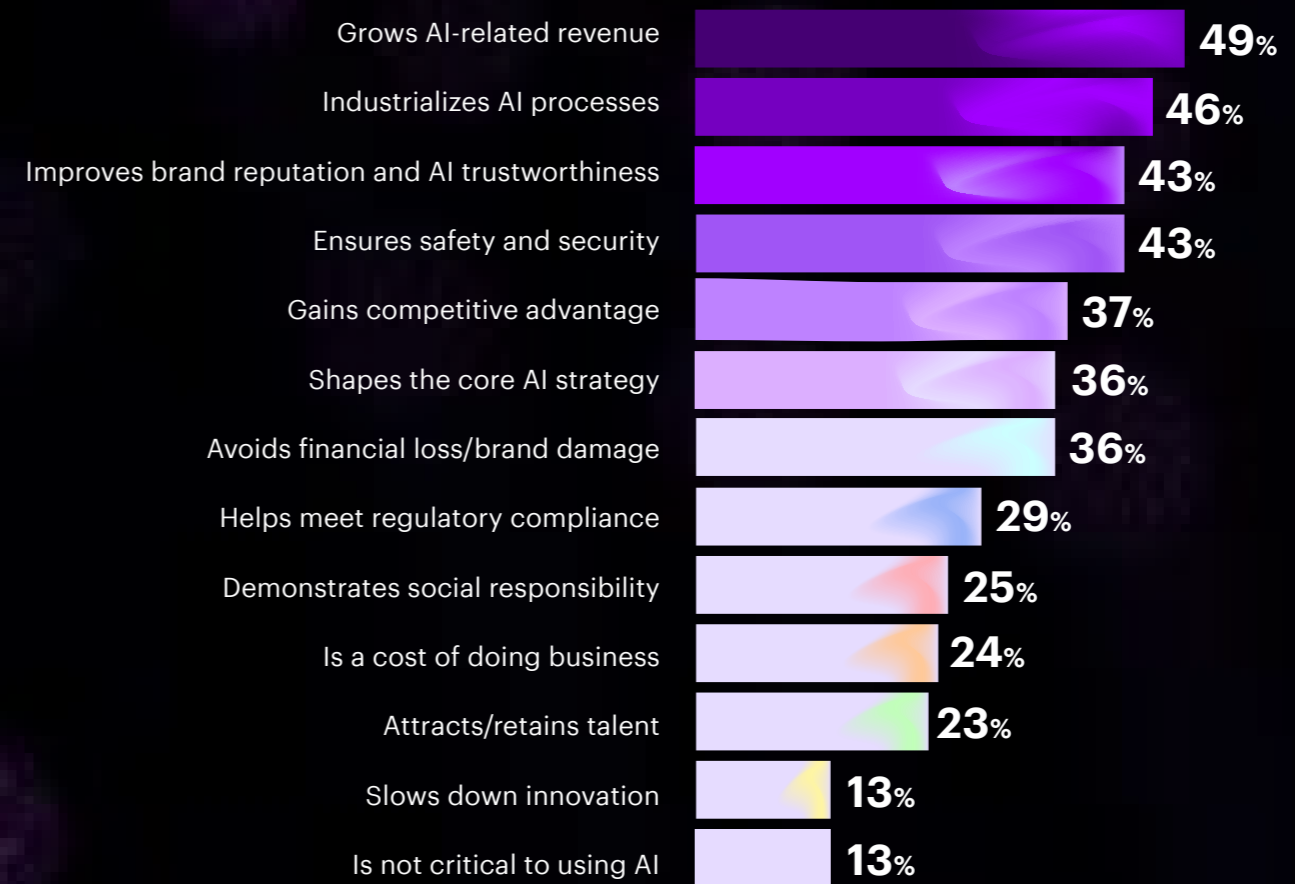
In the past, businesses often made the mistake of treating responsible AI as a mere compliance issue, rather than as an essential contributor to value creation. Fortunately, nearly half of the surveyed firms do not hold this point of view today.

For instance, about half (**49%**) of the companies we surveyed said they view responsible AI as a key contributor to AI-related revenue growth for their firm, while only 29% of companies said that responsible AI is mainly a regulatory and compliance issue (Figure 2). Likewise, 43% of surveyed companies said that responsible AI is an important contributor to protecting their brand's value, while just 24% of firms view responsible AI as simply a "cost of doing business".

49%

of the companies we surveyed said they view responsible AI as a key contributor to AI-related revenue growth for their firm.

Figure 2: C-suite perceptions of responsible AI



Source: Accenture Stanford Executive Survey, N=1,000. See "About the research" for details.



To get a better sense of just how much financial value responsible AI can unlock, we asked executives to provide their own estimates. Their answers suggest that most companies today view responsible AI as hugely important. For instance, **our survey respondents predicted that when a company becomes a pioneer in responsible AI, its AI-related revenue will increase by 18%, on average.** (AI-related revenue is the total revenue generated by AI-enabled products and services.)

The opposite is also true. When responsible AI is not well-established in a company, brand value can be quickly destroyed with its misuse. Our survey respondents, for their part, **estimated that**

a single major, AI-related incident would, on average, erase 24% of the value of their firm's market capitalization. These sentiments further prove why companies are revving up their investments in responsible AI. Of the companies we surveyed, **42% already devote more than 10% of their overall AI budget to responsible AI initiatives; over the next two years, 79% of companies plan to hit this robust spending target.**

A growing number of companies are, in short, prioritizing responsible AI and are spending accordingly. So, how close are they to achieving their responsible AI goals?



Becoming reinvention ready: the milestones of responsible AI maturity

In today's business landscape, continuous change is the new reality. Being set up for continuous change means you need reinvention-readiness in every function and every component of your business. Responsible AI is no different. Reinvention-ready companies have the ability to be agile, they absorb the shifts that are constantly happening in-market and can proactively respond to those shifts at speed, exploiting new opportunities and mitigating unintended consequences. Our research has shown that there are four groups of companies that are at a distinct milestone of their responsible AI evolution, where the goal is to be reinvention ready.

What do we mean by being reinvention ready when it comes to responsible AI? Interestingly, no company has yet reached that milestone, but those who get there first will be responsible AI pioneers. They will have fully operationalized responsible AI efforts as a platform to take a more systemic, future-oriented approach that unlocks the true value of AI.

We've defined four milestones of responsible AI maturity, which ultimately lead to being reinvention ready. Here, maturity is something that will continue to evolve, not a finite state or destination. What is mature today will likely be different in the months to come.

Redefining maturity

To measure companies' responsible AI maturity, we developed a four-stage framework in collaboration with Stanford University. The higher the stage, the greater the progress. Based on our analysis of the responses of the 1,000 executives we surveyed, we then placed organizations at their respective stage, awarding a score for organizational maturity and a separate score for operational maturity.

Note: companies with no responsible AI initiatives were excluded from our analysis.

Responsible AI maturity milestones

Stage 1

Setting responsible AI principles

The company has some foundational capabilities to develop AI systems, but its responsible AI efforts are ad hoc:

- Has defined a set of ethical AI principles and guidelines, including policies and rules for responsible and secure data access and usage
- Has no established processes for governing data quality, data privacy, data security and AI model risk management
- Occasionally conducts risk assessment reviews of data and AI projects
- Has deployed AI project workflows without systemic integration of responsible AI assessments across the data pipeline, model pipeline and AI applications

Stage 2

Establishing a responsible AI program

Following a responsible AI assessment, the organization has put in place the following steps:

- Established a responsible AI strategy for the organization with a well-defined operating model and data and AI governance measures for translating vision into action
- Defined a robust approach and process for AI risk assessment across the data pipeline, model pipeline and AI applications
- Established processes for creating transparency and auditability of training data, model inputs and outputs—with appropriate decision-making
- Designed a framework for monitoring and controls across the data and AI pipeline that can be executed during project workflows
- Implemented processes that are still at an early stage and without a more systemic enablement with tools and technology

Stage 3

Putting responsible AI into practice

The company has systematically implemented the following measures across the organization to help meet the relevant regulatory and legal obligations:

- Operationalized responsible AI strategy with implementation of principles, guidelines and processes through to enablement across the business
- Implemented risk assessment across the data pipeline, AI model and AI applications to enable traceability and transparency across the entire model lifecycle and adherence and compliance through regular audits
- Implemented controls for data sourcing with privacy filtering, anonymization and validation to remove sensitive information and mitigate data bias risks embedded into self-service tooling
- Enabled systemic AI testing with model interpretability tools to ensure explainability, AI model performance testing for bias, accuracy, etc. and help guard that models operate within the required legal, ethical and operational boundaries
- Established a responsible AI control plane with human controls for continuous monitoring across the data and AI value chain to alert and remediate any unintended risks or breaches
- Rolled out a responsible AI academy for employee training and enablement to drive responsible AI adoption

Stage 4

Becoming a responsible AI pioneer

The company has fully operationalized responsible AI efforts as a platform to take a more systemic, future-orientated approach that unlocks the true value of AI:

- Fully operationalized end-to-end systemic responsible AI effort powered by tech platforms, redesigned processes with the right talent and culture established
- Adopted an anticipatory approach to their responsible AI efforts, deploying dedicated resources, processes, etc. to continuously assess current and future risks
- Proactively adapts their data and AI risk management and control processes as the external technology and regulatory environment expands and evolves
- Continuously refines and advances data governance and management practices, employing predictive analytics and real-time data monitoring to dynamically understand and manage the impact of data on AI systems
- Shapes new standards, methods and approaches for safe development and use of AI including data privacy preservation, model explainability and bias measurement, adversarial testing and red teaming
- Recognized as a leader in shaping responsible AI practices and actively engages with external stakeholders—including value chain partners, regulatory bodies and affected communities—to ensure participation and inclusive feedback and support forward-looking regulatory compliance
- Proactively engages with third-party AI vendors to foster improvements, building trusted relationships and new collaborative opportunities to manage third-party AI risks effectively



Responsible AI maturity milestones: examples*

Stage 1

Setting responsible AI principles

An Asian manufacturing company wanted to protect against the risks of unchecked AI use.

It established an internal AI policy and checklist to ensure safety, security, fairness, transparency and accountability—but needed help getting implementation right, including training employees.

This company developed a playbook to support the AI ethics review process that cut in half the time required for field operators to run through the checklist. Now, this company can quickly identify AI ethical risks and integrate AI ethical governance into the core of its business—and carry out its vision responsibly.

Stage 2

Establishing a responsible AI program

A multinational consumer healthcare company wanted to define a clear policy and vision to scale responsible AI across the enterprise and standardize processes and ways of building and deploying AI.

The company did not have an inventory of high-risk AI applications and struggled with the absence of dedicated responsible AI roles and decision-making accountabilities.

They worked with Accenture to conduct a global benchmarking and assessment against the regulatory landscape and draft AI principles and policies, as well as a proposed responsible AI operating model. Risk screening was conducted across the company's AI applications. A risk assessment for higher-risk cases ensures their applications align with principles and regulatory requirements. They also worked with third-party legal counsel to provide a framework for legislation monitoring.

The company is now confident in its AI usage, including its risk management and accountability, to the extent that it plans to publish an external position paper on responsible AI.

Stage 3

Putting responsible AI into practice

Over the past few years, Accenture has undergone its own efforts to increase responsible AI maturity, building on its existing AI principles. The company developed a rigorous responsible AI program and put it into practice. The key elements of the robust program are:


- **Leadership oversight:** Accenture appointed a Chief Responsible AI Officer to oversee the internal responsible AI compliance program, with the Accenture CEO and General Counsel as sponsors. The company's Chief Technology Officer, Chief Responsible AI Officer⁶, General Counsel, and Data and AI Lead oversee the related steering committee. There is also regular board reporting of progress to an audit committee.
- **Principles and governance:** Accenture's approach to developing and deploying AI solutions is founded on a set of principles⁷ that are applied to its own operations as well as its collaborations with clients, partners and suppliers. The company appointed a cross-functional team (including Legal, Security, CIO, Procurement, HR and responsible AI experts, among others) to design and lead the new responsible AI compliance program with the responsible AI principles acting as the North Star and anchor for the design.

A governance framework that implements key principles, policies and standards supports cross-use case supervision.

- **Risk assessment and mitigation:** Accenture takes a risk-based approach in accordance with the EU AI Act and other key frameworks. This meant spending significant time up-front defining the higher-risk use cases. Over the last year, Accenture screened thousands of AI use cases across the company and completed detailed risk assessment and mitigations.
- **Testing and enablement:** The company also spent time on systemic enablement—designing its AI policies, standards and controls and embedding them into technologies, processes and systems. This required building responsible AI technical capabilities, expertise, tools and techniques and developing benchmark testing tools.
- **Talent:** Accenture is also growing its responsible AI skills and talent, with mandatory ethics and compliance training for those Accenture people who are most directly involved with AI, as well as new ethics and AI training through Accenture Technology Quotient (TQ) courses for its 774,000 people.


* The companies in these examples have achieved major elements of the respective stage. The North Star for responsible AI maturity is to become a pioneer. No companies are yet at this stage.



A satellite is shown in space, emitting a beam of light. The background is a dark field filled with a grid of small, glowing light points, creating a sense of depth and technology. The overall color palette is dark with vibrant highlights.

Responsibility reality check: how ready are companies for responsible AI?



A man with light brown hair, wearing a dark blue button-down shirt, is shown in profile from the chest up. He is holding a gold-colored pen in his right hand and looking off to the right with a slight smile. The background is a blurred office environment with a window on the right side.

While good progress is being made, companies are likely experiencing a perception gap between their intention and execution, as even those who perceive themselves to be mature have a long way to go. Our findings indicate that companies may still be underestimating the number of risks they are exposed to, the quantity of measures required and the completeness of how they are implemented. Without a robust set of risk mitigation measures, organizations are not just exposed to existing risks, they will also struggle to adapt to new regulations, anticipate new emerging risks and scale AI opportunities confidently.

Operational versus organizational maturity

To advance responsible AI maturity, companies must translate organizational maturity into a comprehensive set of mitigation measures, or operational maturity. We observed that while organizational maturity has continued to grow over the last two years, there is a significant disconnect with operational maturity.

Operational maturity

Operational maturity is currently a big weakness for most companies. Operational maturity measures the extent to which a company has adopted responsible AI measures to mitigate specific AI-related risks, including those related to privacy and data governance, diversity and discrimination, reliability, security, human interaction, accountability and environmental impact. A high score on operational maturity indicates that a company excels at implementing AI responsibly across all risk areas that apply to their business. Alas, only a small minority of companies appear to be implementing responsible AI with success. We found that **6% of the companies we surveyed have reached the practice milestone of operational maturity and less than 1% of companies are at the pioneer stage** (Figure 3).

Responsible generative AI

When it comes to generative AI specifically, we analyzed firms' ability to mitigate risks for both developers (some companies) and users of the technology (nearly all companies).

For users of generative AI, we assessed mitigation measures applied at each stage of the AI lifecycle i.e. provider selection, evaluation, infrastructure, application, end-user support and monitoring, control and observability. For developers, we focused on AI infrastructure, model development and evaluation, application development, end-user measures and post-deployment monitoring, control and observability. We found that operational maturity for generative AI and for AI in general were similarly low (Figure 3)—**just 13% of companies are at either the practice stage (10% of companies) or the pioneer stage (3% of companies)**.



Organizational maturity

Organizational maturity reflects the extent and effectiveness of an organization's responsible AI processes and practices—as measured by C-suite sponsorship, governance, data and AI risk identification and management, model development, procurement, monitoring and control, cybersecurity and training. A high score on organizational maturity indicates that a company is well-prepared to use AI responsibly.

For many of the companies we surveyed, organizational maturity is a relative strength (note: scores reflect companies' self-reported capabilities). As Figure 3 shows, 72% are either at the practice or pioneer stage (63% and 9% respectively). Only 3%, meanwhile, are at the principles stage and 25% are at the program stage. Compared to two years ago, C-suite sponsorship of responsible AI has increased from 50% to 79%, while companies with a fully-operationalized risk management framework are up from 48% to 69%. This shows that companies are starting to take responsible AI seriously and are adopting a top-down, organization-wide approach, which is a foundational step in reaching maturity milestones.

Prioritizing responsible AI initiatives has a direct impact on maturity

Our analysis shows us that investing in responsible AI positively correlates to higher maturity in all three dimensions: operational, generative and organizational responsible AI maturity. Companies that spend more than 10% of their AI budget on responsible AI are two times more likely to be at the practice and pioneer stages of operational maturity and three times more likely to be at the practice and pioneer stages of responsible generative AI maturity.



Overall responsible AI maturity

To fully understand the current responsible AI maturity landscape, we designed a framework that uses a composite of individual scores across organizational maturity, operational maturity and generative AI maturity.

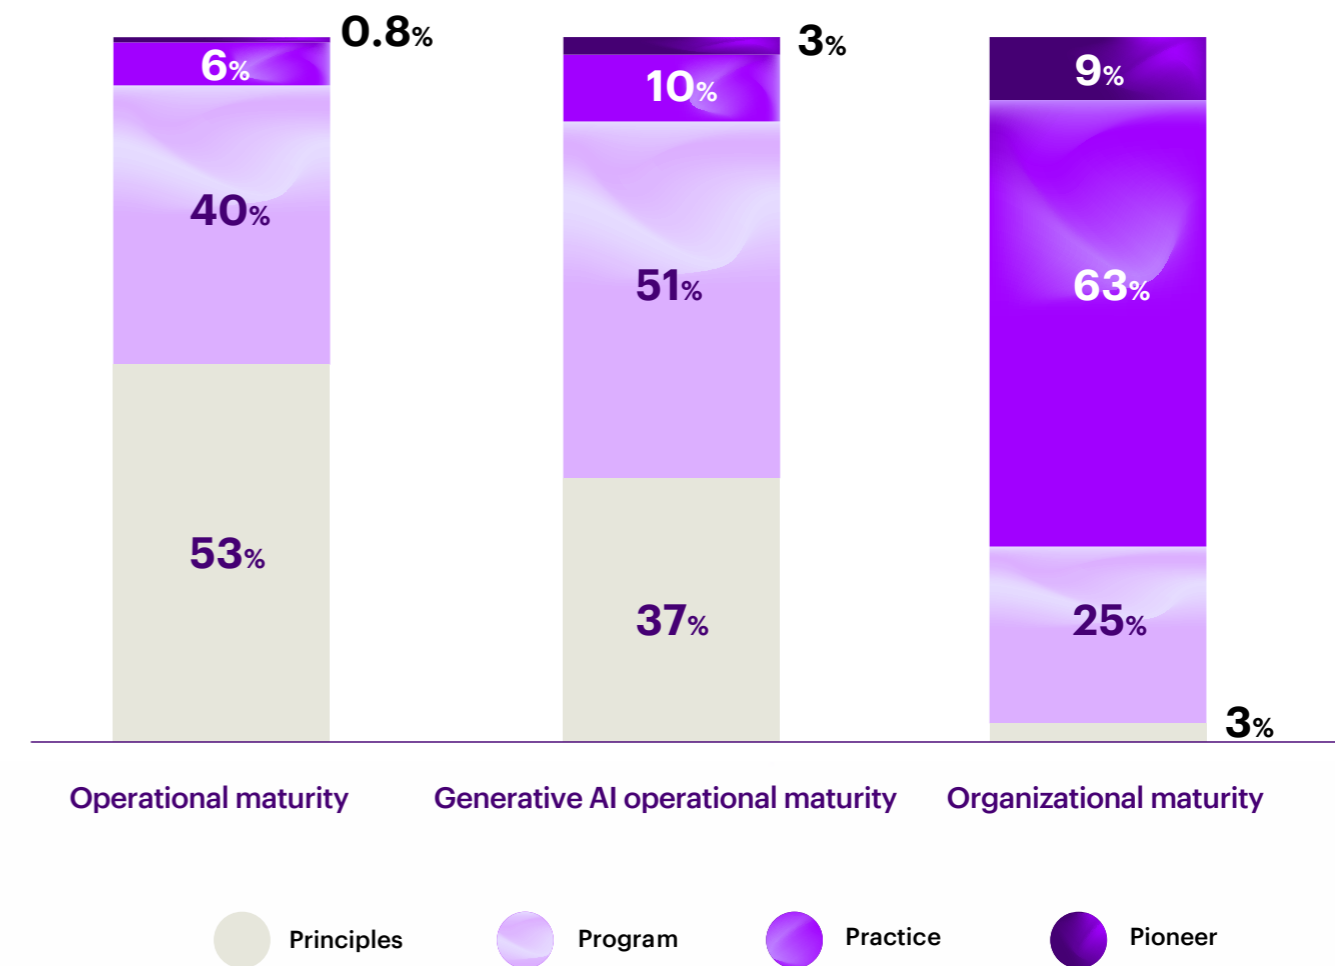
Our research reveals that a vast majority (78%) of companies have established a responsible AI program*. A smaller portion, 14%, have put responsible AI into practice, while 8% are just beginning their journey by setting responsible AI principles. Notably, none of the companies have become a pioneer.

However, this aggregate view obscures the substantial differences that exist between organizational and operational maturity (Figure 3). For example, although at an organizational and operational level pioneers exist, when you combine scoring, you are left with no pioneers. Organizations need to make sure they prioritize organizational maturity and operational maturity. Organizational maturity is the natural place to start and therefore companies are often more advanced on that front. But they must remember that taking a systemic approach to operational maturity is also critically important (as that's when you're really putting the measures into practice in your organization). Otherwise, a perception gap can emerge which leaves companies exposed and unable to adapt to the changing risk landscape.

Our analysis found that there are currently no significant differences in responsible AI maturity across different industries, suggesting a consistent approach to responsible AI across sectors. However, we observed that companies in Asia, for example in India and Singapore, are leading in responsible AI maturity, demonstrating more advanced practices compared to their global peers. Among European countries, Germany stands out, with its companies showing higher levels of maturity in responsible AI compared to other nations in the region. In many cases these are countries that have established broader AI policy, investment and regulation, which likely helps to drive local adoption.

* Companies with no responsible AI initiatives were excluded from our analysis.

Figure 3: Percentage distribution of companies across four levels of responsible AI maturity by operational, generative AI and organizational areas



Source: Accenture Stanford Executive Survey, N=1,000. See "About the research" for details.



What good looks like: the mark of a responsible AI pioneer

When it comes to responsible AI maturity, reaching the pioneer stage should be the new goal. **As noted, a small minority (9%) of the companies we surveyed are responsible AI pioneers for organizational maturity; and even fewer (less than 1%) for operational maturity.** When combining operational and organizational maturity, no company has become a pioneer.

From a geographic perspective, Asia leads globally with the highest number of companies at the pioneer stage, driven by strong performances in Singapore, Japan and India. Asia makes up 34% of organizational pioneers and 37% of operational pioneers. In Europe, Germany stands out as the regional leader, especially when it comes to organizational maturity. This is consistent with what we'd expect, given these countries either have or are exploring regulation at a government level. In North America, where regulation is emerging ad hoc at the individual state,

city or industry level, we see they lag behind Asia and even some emerging regions in operational maturity.

From an industry perspective, Communications, Media and Technology (particularly in High Tech and Software & Platforms) is most advanced in terms of maturity, with 27% of companies advancing in operational maturity and 20% in organizational maturity. Within the Financial Services sector, Insurance and Banking are the most mature.

These findings aren't surprising, when you consider that the Financial Services industry is highly regulated, and Software & Platform players by and large see self-regulation as critical to their future. They've already had to put processes in place to monitor and mitigate the AI they're using. Finally, the Retail and Industrial Equipment industries show strong maturity while Consumer Goods & Services lags.

9%

of companies have reached the pioneer milestone for organizational maturity, and this drops to less than 1% for operational maturity.

When combining organizational and operational maturity, no company is a pioneer yet.



What sets responsible AI pioneers apart?

As we move into the generative AI era, organizations that want to leverage responsible AI as a value lever must build on their organizational and operational maturity. Adopting a future-looking mindset to continuously improve and evolve in line with new technology advances and regulatory changes will be imperative. Responsible AI pioneers are:

Anticipators

To be skilled anticipators, pioneers continuously adapt their risk monitoring and control processes as the external technology and regulatory environment evolves. They systematically align their AI activities with the organization's overall strategic objectives and tolerance for risk. And they have dedicated teams and built-in mechanisms to continuously assess and manage current and future risks, creating a platform for rapid adoption of new technology advances.

When other companies hesitate to deploy the latest advances in AI—such as “agentic” AI (AI systems that exhibit increasing levels of autonomy and the ability to take actions independently)—for fear of the risks, pioneers can move ahead with confidence, knowing that their responsible AI efforts are ready for the challenge.

Responsible by design

Pioneers are “responsible by design”, placing responsible AI at the center of their overarching AI strategy. They understand that responsible AI is crucial to day-to-day operations, providing a robust cross-functional, organizational framework that facilitates decision-making in real-time, supporting the adoption and scaling of new technologies confidently.

Pioneers are considered leaders in AI innovation, investing in future planning and horizon scanning to maintain their competitive edge by perpetually refining and improving their responsible AI governance structure, principles, policies and standards.

Proactive partners

Pioneers also collaborate closely with their partners to manage risks, support multi-stakeholder participation, feedback and inclusive decision-making and support forward-looking regulatory compliance. For example, pioneers organize and lead candid discussions between stakeholders in their broader AI ecosystem. The goal of such gatherings is to build bonds of trust that facilitate the sharing of best practices around responsible AI, while also providing a channel for offering technical support to partners that need it.

By collaborating with partners on responsible AI, pioneers establish mutual trust, strengthening their own maturity and supporting new business opportunities.



Ready, set, grow: five priorities for responsible AI

As our data shows, organizations have reached very different milestones on their responsible AI maturity journey. A company at the principles stage of responsible AI maturity will inevitably have different priorities as it seeks to improve its responsible AI capabilities, than a company at the program or practice stages. Nevertheless, our research and work advising clients has shown that all companies can benefit from focusing on these five priorities to improve their maturity and begin to reap the benefits of AI.

01 Establish AI governance and principles	02 Conduct AI risk assessments	03 Systemic enablement for responsible AI testing
04 Ongoing monitoring and compliance	05 Workforce impact, sustainability, privacy, security	

Priority 1

Establish AI governance and principles

Developing a responsible AI strategy and roadmap that includes clear policies, guidelines and controls is critical for the successful implementation, operationalization and governance of responsible AI practices.

A foundational step in this process is to define, adopt and enforce a set of responsible AI principles (specific to your company priorities, ethics and values) to ensure the ethical design, deployment and usage of AI. Our research shows that 70% of companies have established responsible AI principles, with 54% having translated these into policies.


Implementing a robust AI governance operating model on which responsible

capabilities and controls can be built is another important step.

The good news is that 76% of companies have fully operationalized their governance model—a dramatic increase from just 31% two years ago. This is reflective of companies taking a top-down, cross-functional approach to responsible AI, rather than attempting to manage risks in an ad-hoc manner.

Companies can take this a step further and reinforce this operating model by providing employee training and change management initiatives to promote and sustain responsible AI across the organization.





Priority 2

Conduct AI risk assessments

Understanding risk exposure from an organization’s use of AI is a key component of operationalizing responsible AI. If you don’t know where you’re exposed, how can you protect yourself?

Most companies we surveyed appear to be underestimating the number of AI-related risks they face. For example, when we presented 13 AI-related risks to our survey respondents and asked which of those risks their company was worried about (they could select as many as they wanted), companies selected just 4.4 risks, on average. This underestimation of AI risk is visible in another finding: **over 50% of companies we surveyed do not have a systematic risk-identification process in place.**

Adopting a systematic approach to the screening and categorization of risks from AI use cases is important when conducting risk assessments. Quantitative (scenario analysis, stress testing and key risk indicators) and qualitative (failure mode-and-effects analysis, root-cause analysis, expert human judgement) tools and assessments can help highlight the risks of an organization’s AI use, including fairness, explainability, transparency, accuracy, safety and human impact. Scaling these and other tools across an AI’s lifecycle and value chain will help companies to better identify and respond to the many AI risks they face.

Priority 3


Systemic enablement for responsible AI testing


When firms comprehensively test and scale responsible AI, they deploy a broad range of risk mitigation measures (we evaluated companies' maturity on 44 measures in our survey) across both the AI lifecycle and value chain. Yet **just 19% of surveyed companies had scaled more than half of the risk testing and mitigation measures** that we asked them about.

To comprehensively test and scale responsible AI across the organization, companies need to first develop a reference architecture to seamlessly

integrate client and other third-party tools and services to evaluate risks across the full AI lifecycle (data, model, application).

With this in place, companies can test, fine-tune, integrate and deploy that architecture across the organization. In addition, they need to provide role-based training (such as for customer-service professionals, HR professionals and managerial roles) to enhance employee skills on the latest responsible AI processes and tools.





Priority 4

Ongoing monitoring and compliance

Establishing a dedicated AI monitoring and compliance function is crucial for ensuring the compliant, ethical and sustainable performance of AI models within an organization. This step is particularly important when it comes to generative AI applications, where there is currently less data and model transparency and a far higher frequency in incidents like hallucinations, bias and IP or copyright breaches occurring. As unintended and extended consequences continue to emerge, and companies deploy models or agents that can take actions and make decisions autonomously, the ability to monitor becomes an imperative.

Despite this, **43% of companies have yet to fully operationalize their monitoring and control processes**, making it the weakest element of organizational maturity. Furthermore, **52% of generative AI users do not yet have any monitoring, control and observability measures in place.**

Getting responsible AI right also requires having employees who are focused full-time on the cause.

To do this, companies should start by opening an office tasked with planning, implementing and managing the organization's responsible AI initiatives. That office should be composed of a cross-functional team from a variety of disciplines, with well-defined roles for personnel and a clearly-delineated hierarchy of accountability. Office personnel in the "control tower", for instance, might be accountable for monitoring risks in real time, while other personnel might be charged with monitoring risks on the horizon.

A dedicated AI monitoring and compliance office also needs to be equipped with advanced tools and methodologies to perform at the highest level. For instance, certain tools today can monitor the AI models themselves, track performance metrics and detect anomalies. Methodologies like drift detection can review data to ensure an AI model remains accurate over time; user feedback integration can make models more responsive to the needs of the people who use them.



Priority 5

Workforce impact, sustainability, privacy, security

For a successful responsible AI compliance program, cross-functional engagement must address the impact on the workforce, sustainability, and privacy and security programs across the enterprise.


Workforce

Talent is a top priority for leaders globally and is among the largest inhibitors of transformation success. Responsible AI is no different. Companies must ensure that their people have the right skills, in the right places, at the right times and know what the expectations are from a responsible AI perspective. A large majority (92%) of the companies we surveyed acknowledged that employees (and end users of AI) have important roles to play in mitigating risk.

However, those same companies reported that developing human interaction capabilities is the single biggest challenge they face to improve their operational maturity, highlighting the need for reskilling in the age of generative AI.

As Accenture's recent work with MIT demonstrates⁸, employees are the first line of defense on risk mitigation, but to be an effective defense, they need support to overcome potential hurdles (such as low awareness of or complacency toward risk). A multi-disciplinary approach that combines ongoing, high-quality training and reskilling with behavioral science, responsible design and technical interventions will become critical in ensuring optimized risk mitigation and overall performance.





Sustainability

Awareness around AI carbon emissions continues to grow as researchers look beyond large language model training to inference and the wider AI lifecycle to understand where potential issues lie and solutions that can be applied. For example, training BLOOM, a single multilingual language model, produced 24.7 tons⁹ of emissions—the equivalent of 15 round-trip flights between New York and London.¹⁰ As a result, Amazon, Google and Microsoft have all announced plans to adopt atomic energy to support growing energy needs.

Our survey results showed that **only 36% of companies have set up carbon reduction strategies at the organizational level.** As new measurement and mitigation tools and techniques emerge over the next 12-24 months, companies must make sure they are set up to adopt these to manage AI carbon emissions and reduce cloud compute costs.

Cybersecurity

To help mitigate cybersecurity risk, the need for a dedicated AI monitoring and compliance office is particularly urgent. For example, **71% of surveyed companies said they had an AI-focused cybersecurity response plan in place. But far fewer (40% of firms) had designated a specific team to implement their response plan in the event of an incident.** Although this space is still nascent, companies cannot afford to take a reactive approach.




Conclusion: Turn AI risk into business value

If the pursuit of responsible AI were ever merely a compliance afterthought, those days are long gone. Companies today know that to maximize their investments in generative AI and other AI technologies, they need to put responsible AI front and center.

This is, of course, easier said than done, for the challenges of scaling responsible AI across the organization are daunting. AI-related risks are exploding. New AI-focused laws and regulations are germinating everywhere. And AI value chains are growing in complexity, especially as more companies become both developers and buyers of AI models.

Companies must embrace the five priorities above and become responsible AI pioneers if they want to stay competitive. As part of these efforts, companies must pursue an anticipatory mindset, commit to continuous improvement and extend their focus beyond their organization to their entire value chain and wider AI ecosystem.

The reward for becoming a responsible AI pioneer will be considerable: consistently turning AI risk into tremendous business value.



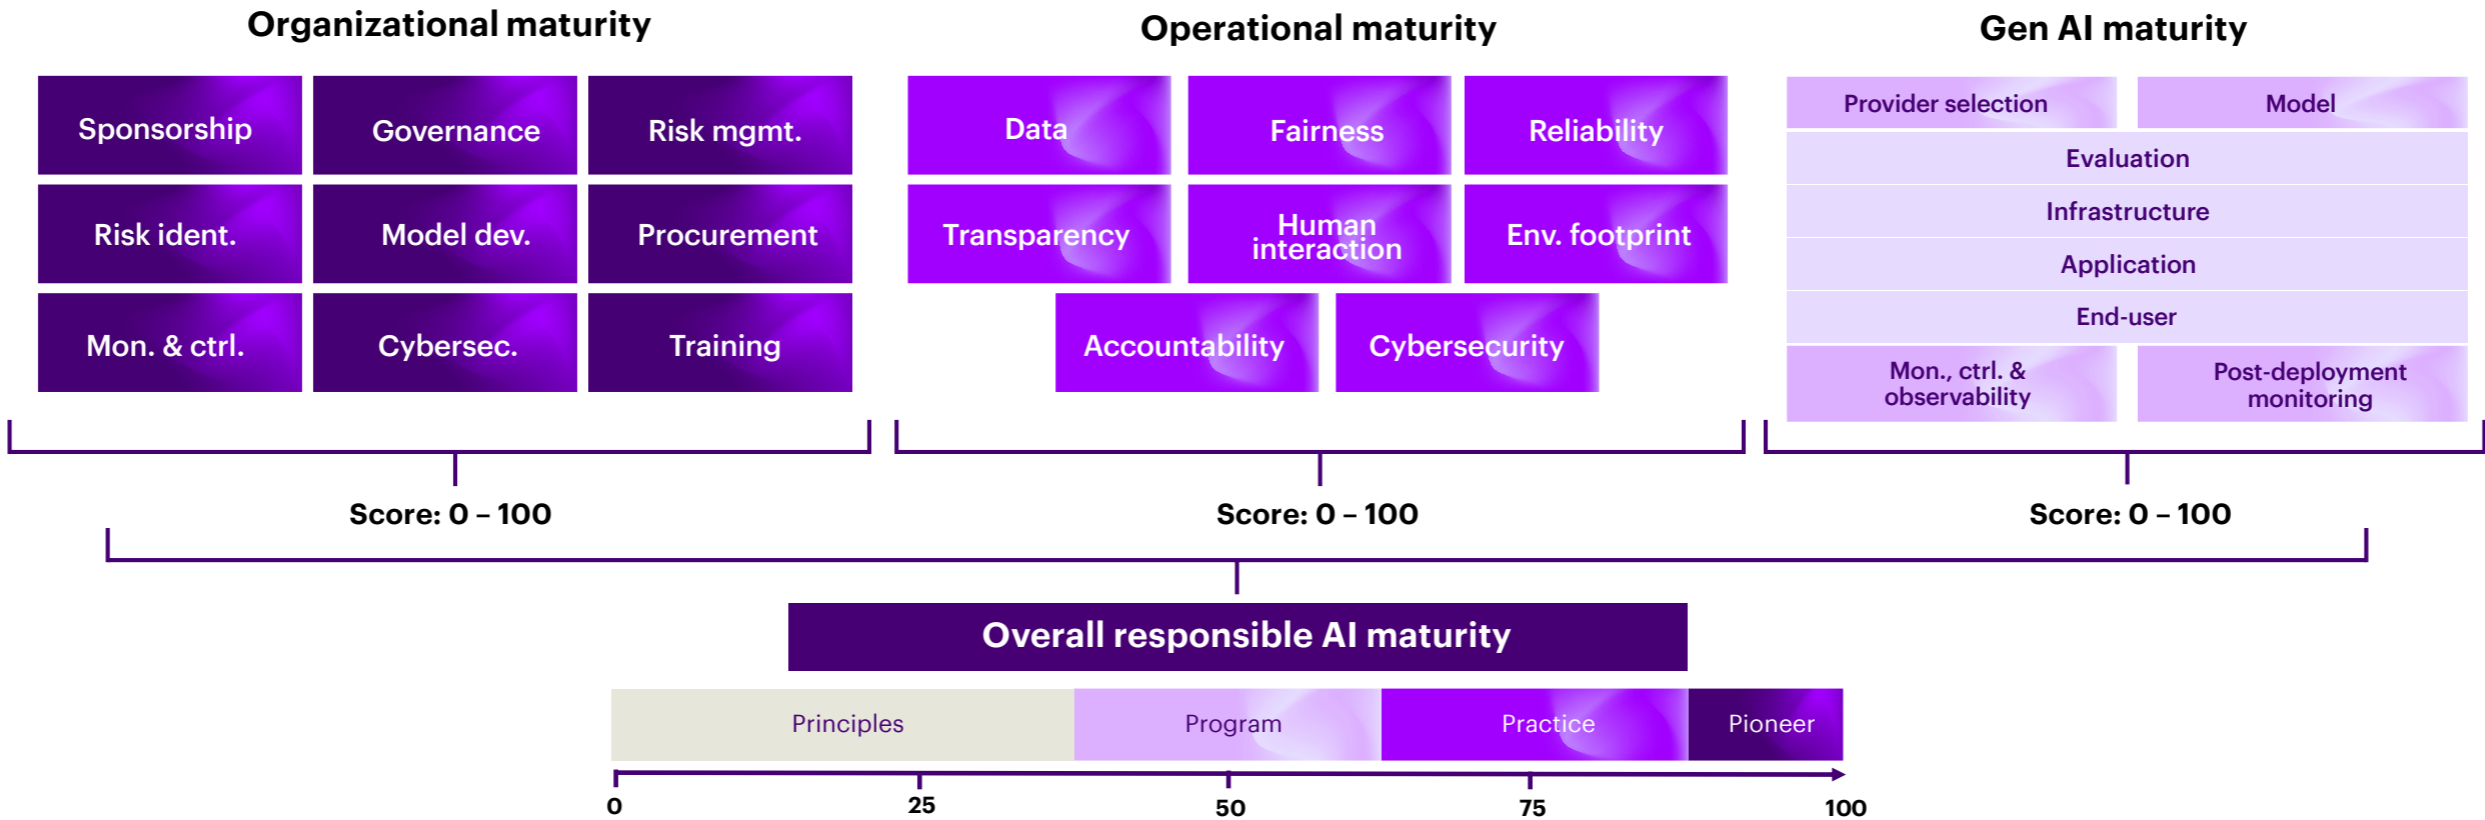
About the research

From January to March 2024, Accenture surveyed 1,000 companies across 19 industries, headquartered in 22 countries. The sample includes a diverse group of business leaders, including CEOs, C-suite executives, Board members and directors. The questionnaire for the survey was co-developed with Stanford University, enabling a robust and comprehensive approach. The methodology of the responsible AI maturity index is based on three core pillars: organizational maturity, operational maturity and generative AI maturity. Each of these pillars is assigned equal weight and their combined scores create the overall responsible AI maturity index, offering a holistic view of companies' readiness in the AI landscape.

Region	Sample
North America	27%
Asia	22%
Europe	30%
Central and South America	9%
Rest of the world	12%
TOTAL	100%

Industry (grouped)	Sample
Aerospace, Automotive & Transport	16%
Communications, Media & Technology	16%
Financial Services	15%
Healthcare & Life Sciences	10%
Products	16%
Public Services	6%
Resources	21%
TOTAL	100%

About the Accenture-Stanford responsible AI maturity scoring model



Authors



Arnab Chakraborty

Chief Responsible AI Officer, Accenture



Karthik Narain

Group Chief Executive—Technology and
Chief Technology Officer, Accenture



Senthil Ramani

Global Lead – Data & AI, Accenture



Acknowledgements

Research director

Patrick Connolly

Research team

Philippe Roussiere

Praveen Tanguturi

Jakub Wiatrak

Shekhar Tewari

Devraj Patil

Dikshita Venkatesh

David Kimble

Marketing team

Sophie Burgess

Erika Marshall

Dries Cuypers

Micaela Soto Acebal

Mark Klinge

References

1. An “AI value chain” incorporates not only the internal processes of a company, but also those of its partners and customers in creating, deploying and maintaining an AI system.
2. **The 19 industries were:** aerospace/defense, automotive, banking, capital markets, chemicals, telecommunications/media/entertainment, consumer goods and services, energy, healthcare, high tech, industrial equipment, insurance, life sciences, natural resources, public services, retail, software/platforms, travel/transport and utilities. **The 22 countries were:** Argentina, Australia, Brazil, Canada, China, Denmark, Germany, Finland, France, India, Italy, Japan, Mexico, Norway, Saudi Arabia, Singapore, South Africa, Spain, Sweden, United Arab Emirates, United Kingdom and United States.
3. <https://www.accenture.com/us-en/insights/data-ai/new-data-essentials>
4. <https://incidentdatabase.ai/>
5. <https://www.comparitech.com/blog/information-security/data-breach-share-price-analysis/>
6. <https://newsroom.accenture.com/news/2024/accenture-takes-new-steps-to-help-clients-scale-generative-ai-responsibly>
7. <https://www.accenture.com/us-en/case-studies/data-ai/blueprint-responsible-ai>
8. <https://sloanreview.mit.edu/article/nudge-users-to-catch-generative-ai-errors/>
9. <https://sloanreview.mit.edu/article/how-developers-can-lower-ais-climate-impact/>
10. <https://www.jmlr.org/papers/volume24/23-0069/23-0069.pdf>

About Accenture

Accenture is a leading global professional services company that helps the world's leading organizations build their digital core, optimize their operations, accelerate revenue growth and enhance services—creating tangible value at speed and scale. We are a talent and innovation-led company with 774,000 people serving clients in more than 120 countries. Technology is at the core of change today, and we are one of the world's leaders in helping drive that change, with strong ecosystem relationships. We combine our strength in technology and leadership in cloud, data and AI with unmatched industry experience, functional expertise and global delivery capability. Our broad range of services, solutions and assets across Strategy & Consulting, Technology, Operations, Industry X and Accenture Song, together with our culture of shared success and commitment to creating 360° value, enable us to help our clients succeed and build trusted, lasting relationships. We measure our success by the 360° value we create for our clients, each other, our shareholders, partners and communities.

Visit us at www.accenture.com.

About Accenture Research

Accenture Research creates thought leadership about the most pressing business issues organizations face. Combining innovative research techniques, such as data-science-led analysis, with a deep understanding of industry and technology, our team of 300 researchers in 20 countries publish hundreds of reports, articles and points of view every year. Our thought-provoking research developed with world leading organizations helps our clients embrace change, create value and deliver on the power of technology and human ingenuity. For more information, visit Accenture Research on www.accenture.com/research.

Disclaimer: The material in this document reflects information available at the point in time at which this document was prepared as indicated by the date in the document properties, however the global situation is rapidly evolving and the position may change. This content is provided for general information purposes only, does not take into account the reader's specific circumstances, and is not intended to be used in place of consultation with our professional advisors. Accenture disclaims, to the fullest extent permitted by applicable law, any and all liability for the accuracy and completeness of the information in this document and for any acts or omissions made based on such information. Accenture does not provide legal, regulatory, audit or tax advice. Readers are responsible for obtaining such advice from their own legal counsel or other licensed professionals. This document refers to marks owned by third parties. All such third-party marks are the property of their respective owners. No sponsorship, endorsement or approval of this content by the owners of such marks is intended, expressed or implied.

Copyright © 2024 Accenture. All rights reserved. Accenture and its logo are registered trademarks of Accenture.