Infosys®
Navigate your next

RESPONSIBLE ENTERPRISE AI IN THE AGENTIC ERA

Infosys® | Knowledge Institute

Contents

Executive summary

Artificial intelligence (AI) is a game-changer, but it also brings risks. Responsible AI (RAI) — the practice of developing and deploying AI systems ethically, safely, and transparently — has emerged as a critical business capability in this context.

The Infosys Knowledge Institute's survey and analysis of 1,500 senior executives involved in enterprise AI governance and execution in North America, western Europe, and Australia and New Zealand (ANZ) uncovers the scale of this risk. This report also identifies many positives about how senior leaders view the need to deploy enterprise AI responsibly. Yet it also reveals large gaps in the ability of many enterprises to safely deliver AI implementations.

AI risks are real and can be severe

An astonishing 95% of C-suite and director-level executives in our survey reported negative consequences of enterprise AI usage in their company in the past two years. These range from privacy violations, inaccurate predictions, to bias or regulatory noncompliance, with almost all involving financial, reputational, or legal damages. More worrying is that almost three-quarters of companies cited damage that was at least considered "substantial," with 39% claiming the damage was "severe" or "extremely severe."

RAI drives business and AI growth

The good news is that this research shows that the best-practice RAI methods and practices can reduce the risk and severity of damages when enterprise AI deviates from expected behavior. Moreover, 78% of senior leaders view RAI as a critical enabler for business growth. Indeed, our research shows that companies that invest in larger RAI teams are able to manage more enterprise AI projects and have a higher volume of successful AI deployments.

Execution and processes are weak

The RAI enthusiasm masks the fact that most companies are not executing it effectively. Only 2% of companies we surveyed met the full standards set in our internal RAI capability benchmark. We termed these high achievers "RAI leaders," with 15% meeting roughly three-quarters of the standards. However, 83% of companies deliver RAI in a piecemeal manner. Those leaders in RAI can expect 39% lower financial losses, 18% lower average severity from their AI incidents, and reduced RAI cost as a proportion of total AI spend.

Most executives blame the lack of resources and rapidly evolving regulations for their weak RAI processes. On average, leaders seek an additional 30% of RAI spending. Yet RAI spending already accounts for 25% of overall AI costs, while financial losses from enterprise AI incidents amount to only 8%. That's a high-risk premium to bear.

Prepare for the agentic AI future

As agentic AI systems — software that makes decisions and acts independently to achieve goals — operate with increasing autonomy, embedded RAI safeguards become business-critical, not optional. Rather than viewing RAI as a bolt-on exercise, these capabilities should be embedded in a platform that engages proactively with the business to identify agentic enterprise AI use cases, supported by an RAI office.

Four specific steps will reduce apprehension about RAI, close gaps, and deliver more benefits. Embedding these practices will not only help companies join the elite 2% of their peers that demonstrate effective RAI, but also ensure they are primed for an era of competition where only the most fail-safe organizations survive.

1. Learn from the leaders.
2. Combine the product and platform operating models.
3. Build RAI guardrails into a platform.
4. Establish a proactive RAI office.

AI risks are widespread

AI and risk go hand in hand — or at least, it seems so in this era of early experimentation. Of the 1,500 senior executives that were surveyed by the Infosys Knowledge Institute in the US, Canada, UK, Germany, France, and ANZ, all were still developing their enterprise AI muscles.

Within this sample, less than a quarter of AI implementations had succeeded in delivering some business value, with nearly 40% being canceled or failing to achieve their objectives. (Note: We discuss which AI use cases derive the most value in our adjacent report, AI Business Value Radar 2025).

Yet almost all executives (95%) have experienced at least one type of problematic incident from their use of enterprise AI. In fact, on average, executives claimed to have experienced around 2.5 different types of AI incidents (Figure 1).

Many of these incidents are similar in type to those that most mature enterprises will be used to defending against outside of their AI work: Privacy violations, security breaches, ethical and discrimination issues, systemic failures, or regulatory noncompliance.

One executive at a large manufacturing conglomerate spoke to us about the risk of data loss using third-party AI tools. Increasing guardrails around these models were critical to safeguard sensitive data, and in some scenarios, building internal models even at high expense was seen as a sound strategy to avoid privacy and ethical violations. But

Figure 1. AI incidents experienced by enterprises

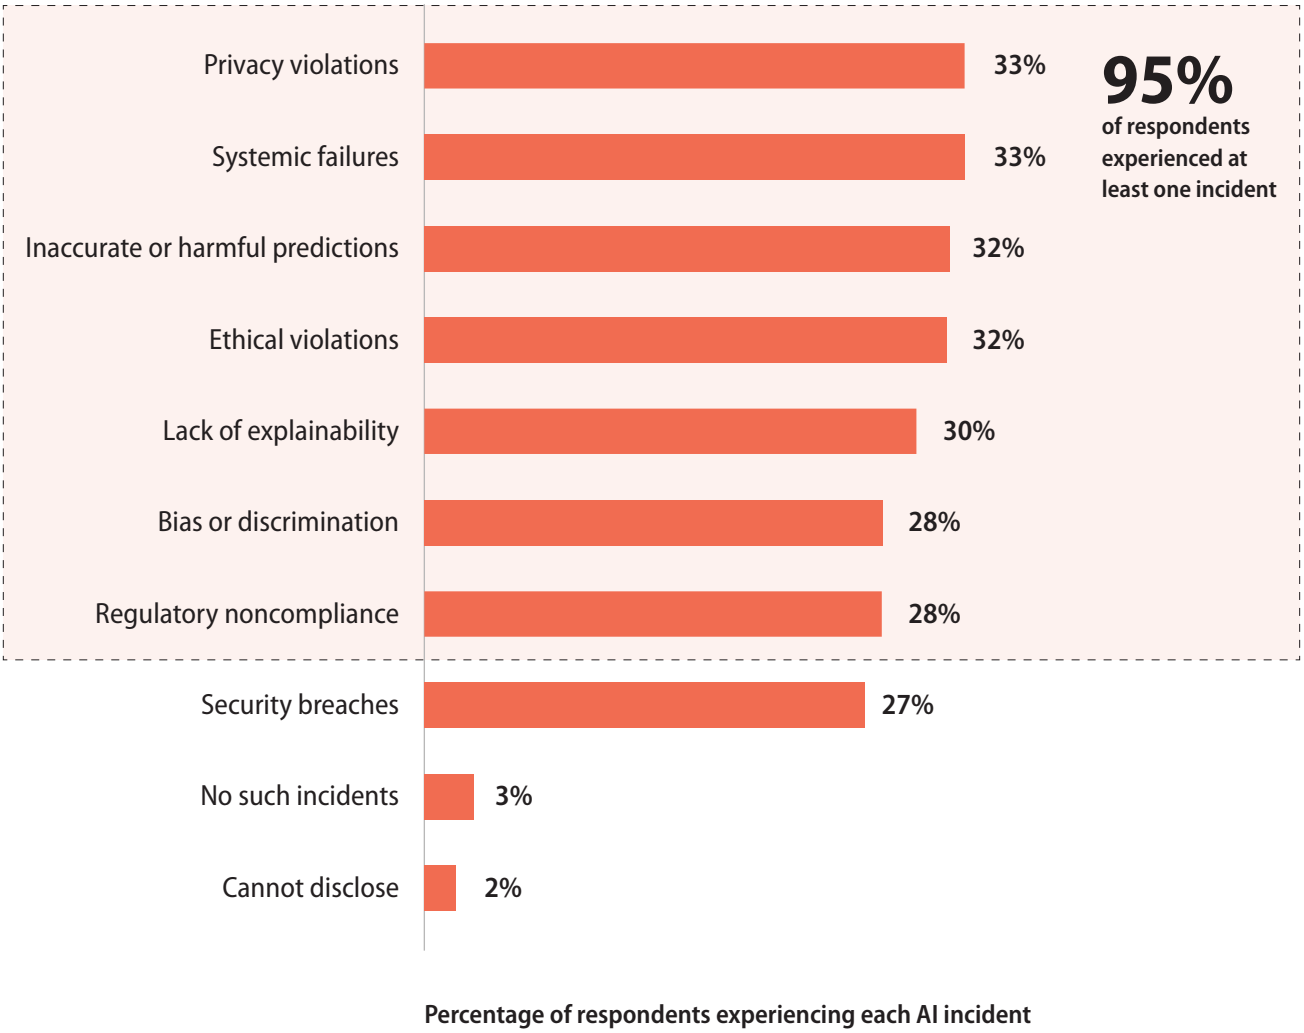

N = 1,502

Source: Infosys

enterprise AI also creates the potential for damage to be caused through inaccuracy or poor explainability of its complex algorithms.

Regardless of the issue type, the scale, speed, and autonomous nature of enterprise AI differentiate these incidents from those that could be caused within more traditional business process problems. AI errors can inflict damage faster and more widely than a simple database error or a rogue employee.

Almost three-quarters (72%) of executives who experienced damage from an enterprise AI deployment rated it at least "moderately severe" — substantial but recoverable damage. But almost 40% rated the damage much higher, as "severe" or "extremely severe," defined as damage that threatened a company's existence (Figure 2).

However, it is worth putting this gloomy data into context. Most enterprise AI initiatives are still in the early stages of development. While respondents commonly rate the damage from problematic incidents as serious, the scope of these deployments is often quite limited — and in many cases so is the actual impact.

Most of the time (77%) the damage incurred from an AI incident is a direct financial loss to the business. This is distinct from financial losses incurred from reputational damage or legal fines, which was not measured by the survey. Such reputational and legal losses are also relevant, as they were experienced by half of our senior executive respondents (Figure 3).

Yet the size of these financial losses, such as

Figure 2. Damage from AI incidents is often severe

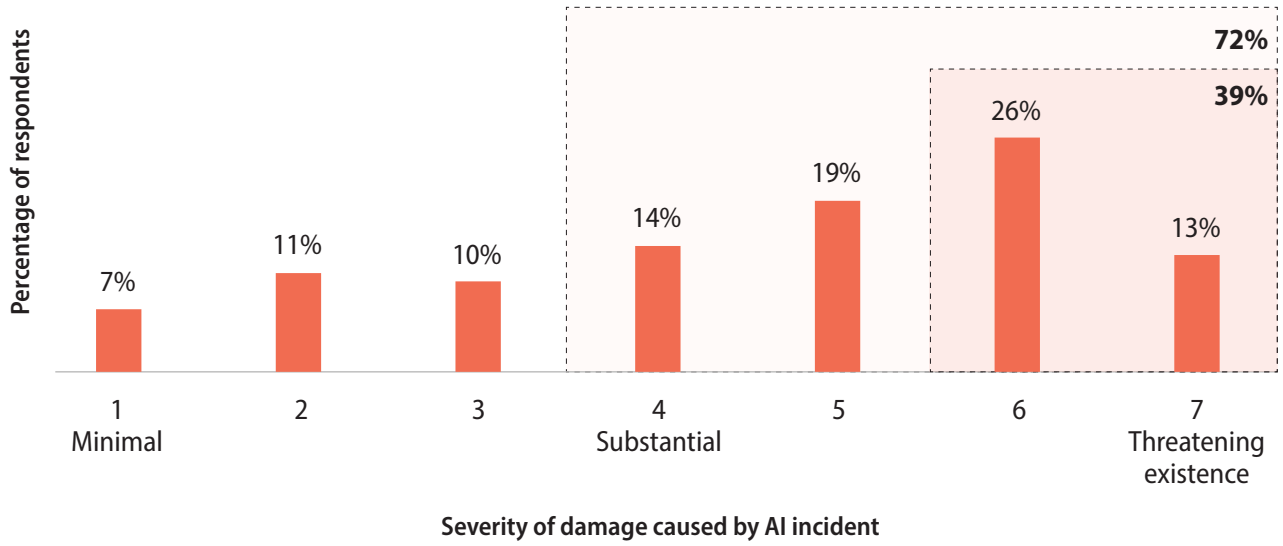

***N* = 1,418**

Source: Infosys

lost revenue or increased costs due to an AI error, is relatively speaking, quite small. The average company in our sample reported financial losses from enterprise AI incidents of about $800,000 over two years (Figure 4). In total this equates to between $750 million and $1.5 billion across the sample, and when extrapolated, represents an annual cost between $1.4 billion and $2.9 billion globally across all businesses ($2.1 billion on average).

It's worth noting that these financial losses were incurred by deployment of enterprise AI algorithms with unintended consequences. They do not cover the costs of privacy, security, or data breaches related to data analytics and management. Losses from such activities are well publicized and can reach huge amounts. Even these sales and cost-related losses may be overshadowed by penalties related to future regulations. As new regulations become established, the potential for fines will increase substantially.

These purely AI-related financial losses within our sample only represent about 8% of companies' total AI spending. Not insignificant, but also not as large as potentially expected, given how frequently this type of loss has been incurred.

Our survey respondents seem to agree. When comparing executive ratings of the damage severity caused by AI incidents, we find the proportion of respondents rating the damage severity as "severe" or "extremely severe" is 12% higher for companies reporting reputational damage than for those without.

This suggests that executives find reputational damage much more threatening to their business than financial losses. Indeed,

Figure 3. Financial losses are the most common result of an AI incident

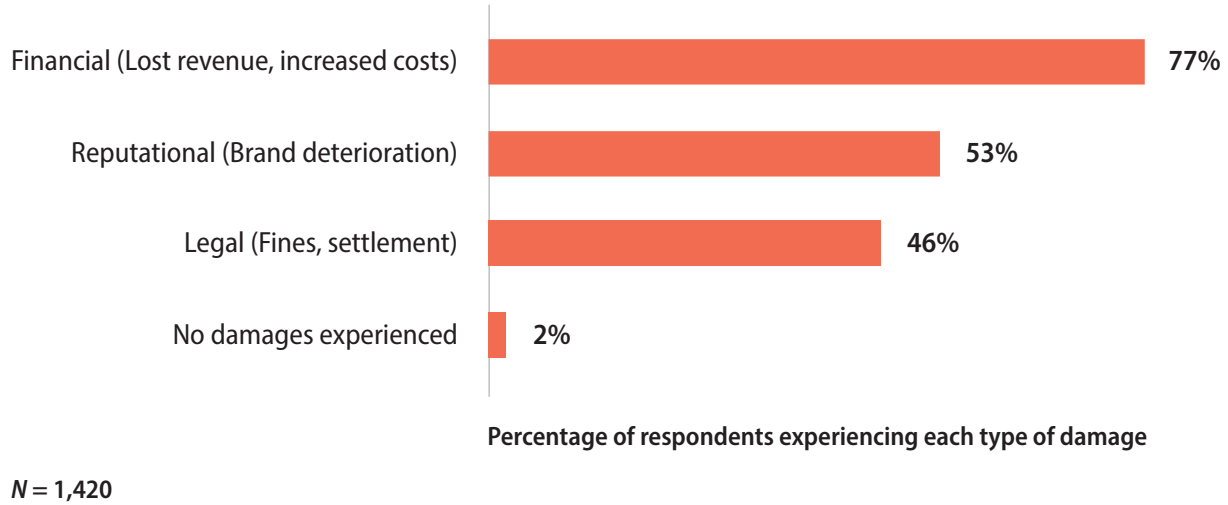

Source: Infosys

our senior executive interviews revealed a certain hierarchy of impact that often began with a legal, regulatory or financial loss, and then extended to a much more damaging reputational loss.

Of course, it matters where enterprise AI is deployed, and what function it performs. If it is used in customer-facing or decision-critical applications, reputational risk increases, with consequences potentially becoming highly visible in a short period. Just look at how discriminatory loan denials or algorithmic misjudgments have become front-page drama — in these cases, the AI system is not a standalone bulkhead, but interwoven and culpable with the culture and values of an organization.

Figure 4. Financial losses from AI are significant but not serious

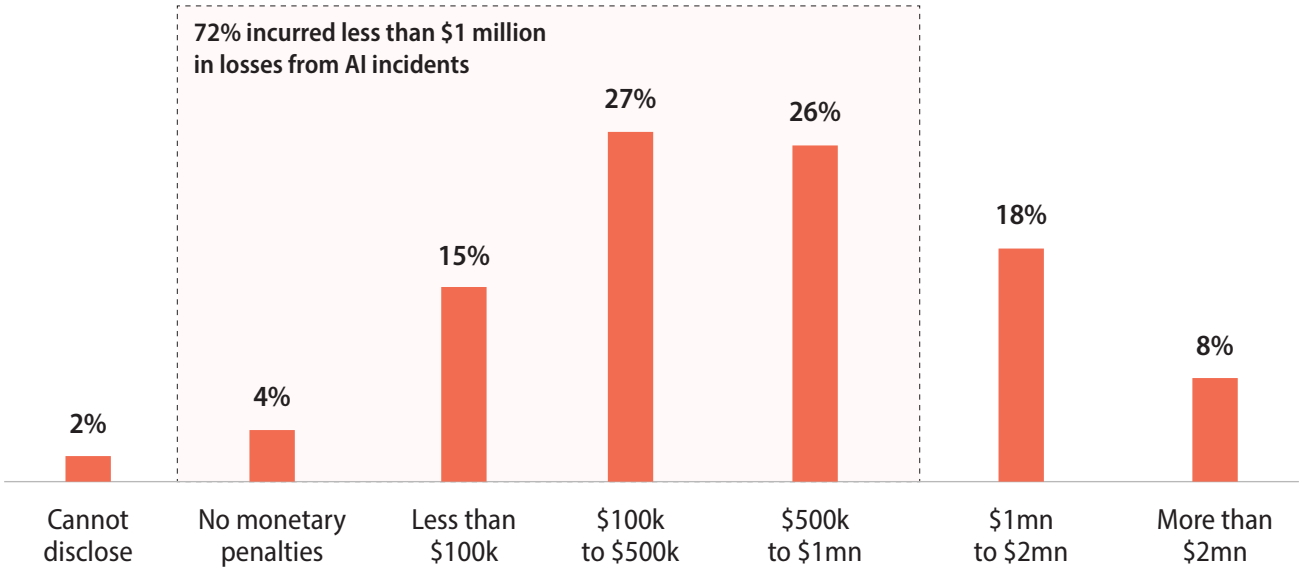

***N* = 1,502**

Source: Infosys

RAI enables growth

A highly positive finding in our research is that RAI practices are perceived to drive business growth. Around 78% of the survey respondents viewed RAI practices as having a positive impact on their business growth, with 15% seeing it as having no impact, and only 7% feeling that it was holding back growth.

Perhaps even more positive is that most executives in our survey welcome new AI regulations, mainly because such regulations will provide clarity, confidence, and trust in enterprise AI both internally and for their customers (Figure 5). These findings allay concern that businesses are rushing into enterprise AI without regard for the risks, or that they see regulations as a blocker to progress.

As one chief experience officer from a large global bank said: "Regulations are a good

thing. It cuts down our work as we don't need to reinvent the wheel ourselves. It's one less thing to worry about and invent. We don't think it will slow us down; rather, regulations will make our solutions stronger."

Our data backs up the executive perception that RAI drives growth — or at least, enterprise AI growth. We found that the larger the RAI team, the more enterprise AI initiatives a company can undertake at a given time. Companies with RAI teams greater than 25 full-time employees can expect to have worked on over 100 enterprise AI initiatives in the past two years.

This represents almost 24% more enterprise AI initiatives than companies with RAI teams of between five and 25 members, and 1.5 times the number of AI initiatives with RAI teams with five or fewer members. It seems logical: The bigger the RAI team, the more enterprise AI initiatives they can approve for testing and deployment, and the more safeguards are in place as growth continues.

There can be other benefits to having a larger RAI team. This is particularly relevant for companies involved in multiple geographies, working in a highly regulated field, or adopted standards such as ISO 42001 (a

Figure 5. Increased clarity and confidence from rules will drive enterprise AI growth

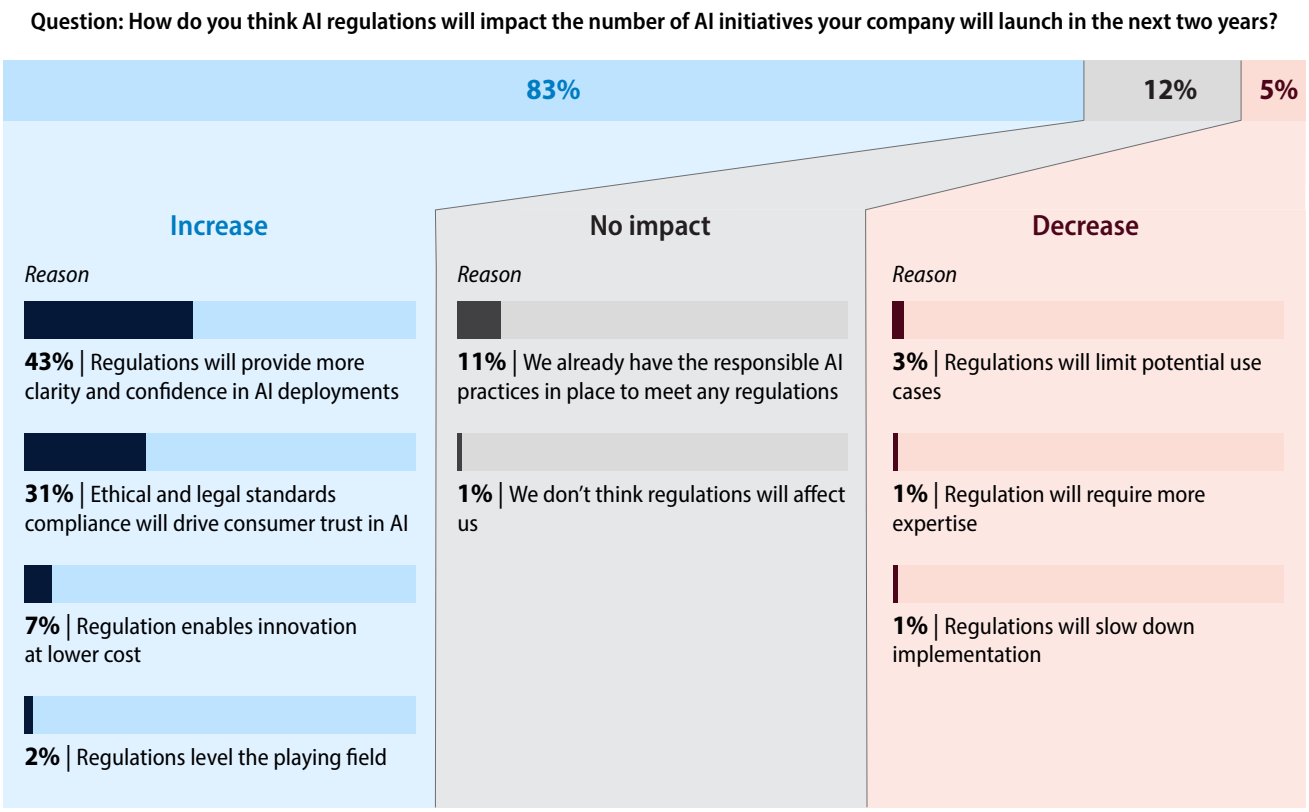

N = 1,502

Source: Infosys

framework to establish, implement, maintain, and continually improve RAI practices), OWASP, or the NIST RMF.

Indeed, our data shows that larger RAI teams correlate with larger companies — suggesting that the more business areas and the larger the complexity of an organization, the more RAI team members are required. Our data also shows that larger RAI teams correlate with more AI initiatives, and this therefore increases the nominal amount of successful AI initiatives (Figure 6). As one executive from the retail industry told us: "We don't differentiate between RAI and AI. We need to apply AI with the right ethics, with the right governance, and with the right security — built in from day one."

However, size alone does not necessarily correlate with success in enterprise AI. The success rate (successful deployments as a proportion of total initiatives) gets significantly lower as team size grows — falling from 24% to 21%. Also worth noting is that companies with larger RAI teams are not always able to deliver better quality RAI outcomes. In our sample, companies with RAI teams of more than 25 members on average suffer 16% higher financial losses from rogue enterprise AI implementations than those with smaller teams. And the severity score of these incidents tends to increase as RAI teams get larger as well.

These less attractive outcomes are not directly due to the size of the RAI team itself,

Figure 6. Larger RAI teams deliver more AI success, but inefficiently

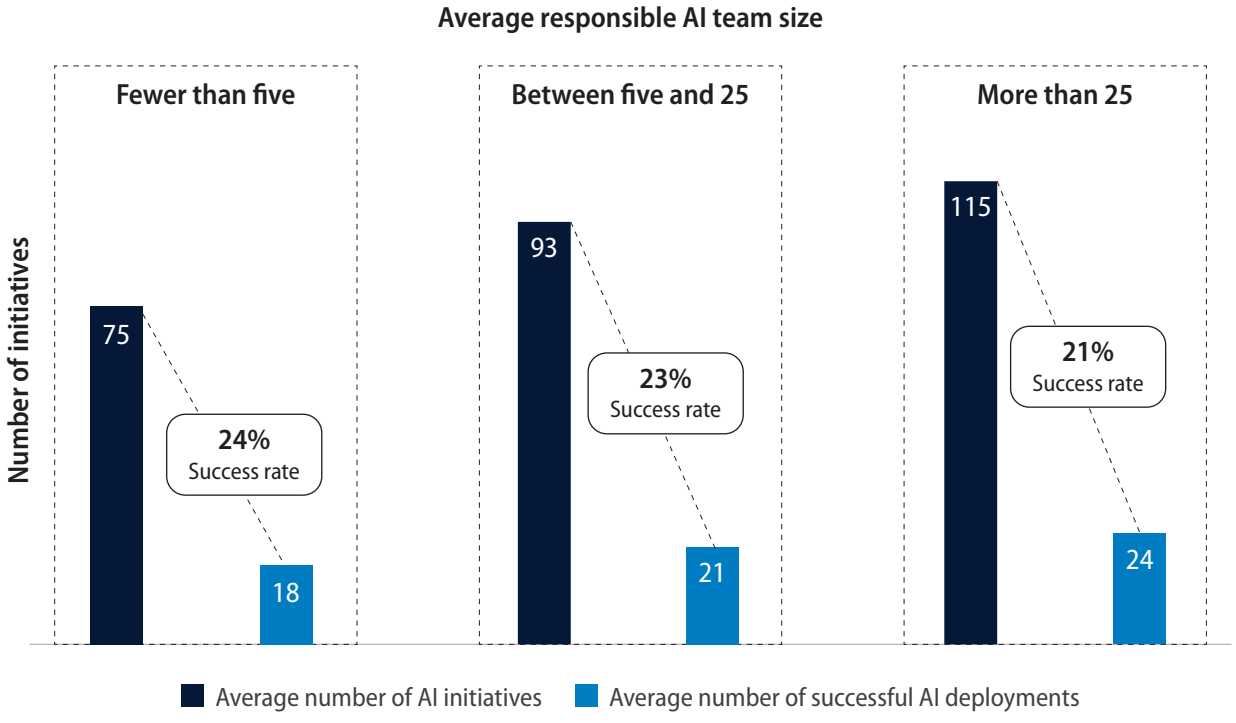

***N* = 1,029**

Source: Infosys

however. More likely, the larger volume of AI initiatives enabled by a larger RAI team means that a larger proportion of these will be less successful. After all, not all AI initiatives will succeed, and it's not the responsibility of most RAI teams to judge their success. Also, a larger RAI team is able to detect and track a larger number of risks and incidents, learning faster and expanding their reach further, which naturally should increase the reporting of these issues.

Regardless of causality, it's clear that currently most companies' RAI functions are effective at enabling AI experimentation that is safe and low risk, but that they are not necessarily contributing to the value that such AI initiatives can bring to the business. In other words, RAI may ensure safety, but it is not yet linked to business success.

This is something RAI leaders should bear in mind when pitching for more resources, as we expect most will be in the coming year. On average, RAI leaders believe their funding should be increased by 30% to deliver RAI effectively (Figure 7). But this money should not be spent just on making RAI bigger. It also needs to be focused on making RAI better. And being better means being safer and more efficient, while supporting better outcomes from AI for the business.

Figure 7. Organizations feel they are underinvesting in RAI by around 30%

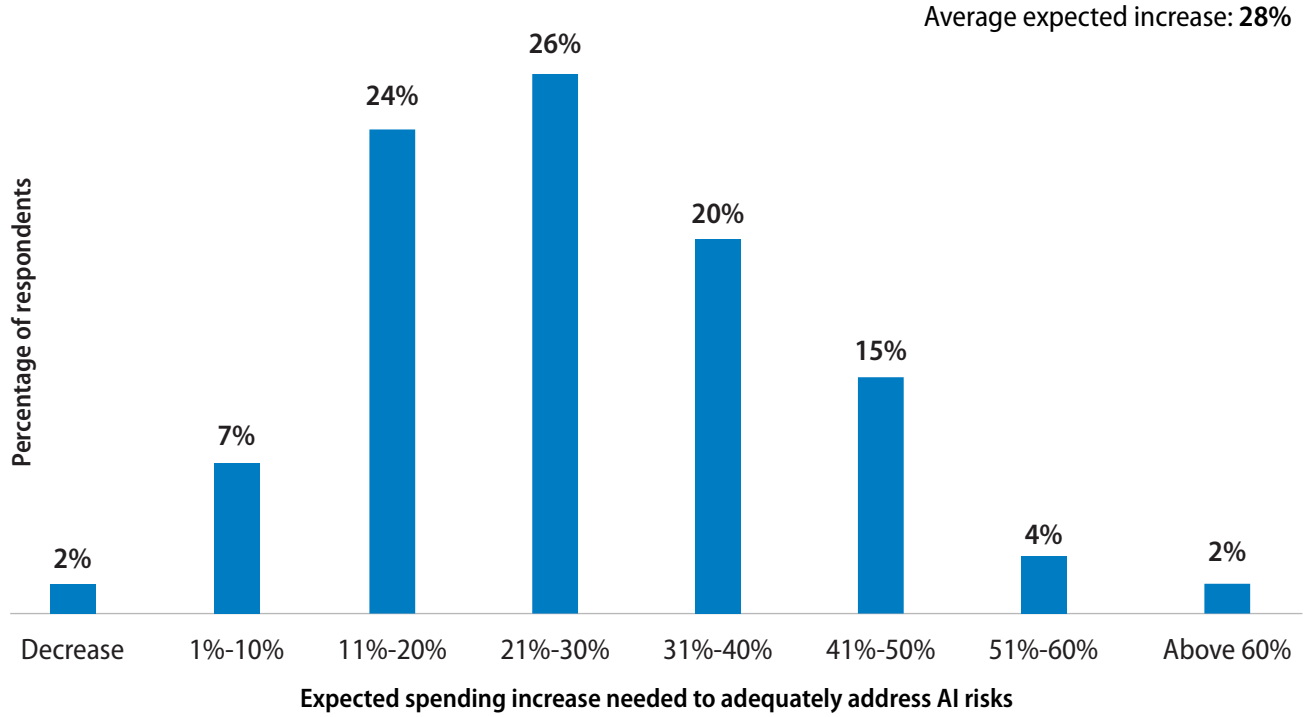

N = 1,502

Source: Infosys

Optimizing RAI

Infosys has its own set of best practices for RAI, covering the dimensions of trust, risk mitigation, data and AI governance, and sustainability. We call this the **RAISE BAR** (**R**esponsible **AI S**tandards **E**valuation: **B**enchmark **a**nd **R**eadiness). This research analyzed companies' current RAI procedures based on these standards and found those that have implemented these procedures experience a lower severity of enterprise AI risk, at a lower cost.

Our research scored respondents on each dimension of the RAISE BAR. Those respondents that met these standards across all four dimensions were termed RAI leaders. Those in the next category, who had achieved roughly three-quarters of the standards, were termed RAI followers. Then we have the RAI beginners, who have only implemented a few of the standards, and finally the RAI laggards, who have barely implemented any.

Figure 8. Leaders' incident costs reduced

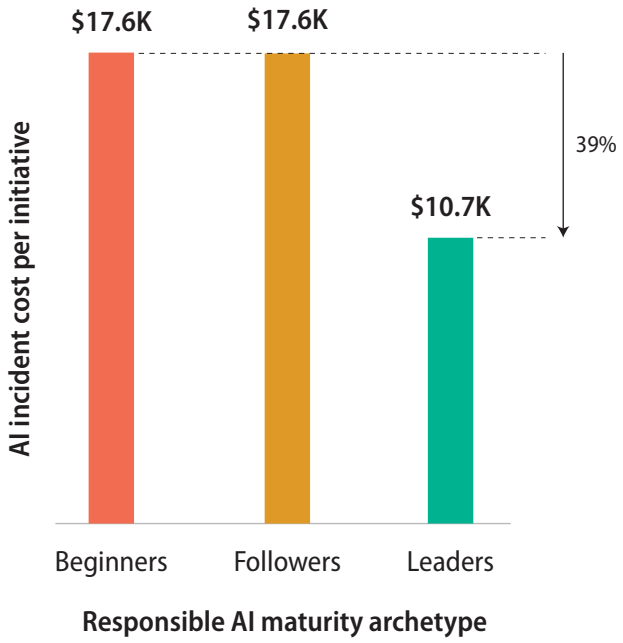

***N* = 1,029**

Source: Infosys

The results are enlightening. RAI leaders experience significantly lower financial costs (lost revenue or increased costs) per AI incident. Indeed, they average 39% lower costs than followers or beginners (Figure 8). Leaders also spend less on RAI as a proportion of AI spend — 21% compared to 25% for followers and beginners — while they experience 18% lower severity of AI incidents (Figure 9).

Another positive aspect of doing RAI well is evidenced by the fact that the companies with the highest RAISE BAR score also experience fewer types of AI incidents (Figure 1). That said, we see in the data that on average, those with higher scores have experienced five or more AI incidents, as compared to those with lower scores that experience three or fewer types of AI incidents.

This suggests companies that mature their RAI capability are doing so by experiencing AI incident types, and as they mature, they identify more AI incident types. But when they complete their RAI capability development and achieve a high standard, they manage to reduce the number and range of AI incidents that occur in their business. Some leading companies are practicing RAI reactively, learning as they go; but there are also some RAI leaders practicing RAI proactively, effectively avoiding further AI incidents.

Figure 9. Leaders see lower costs, severity

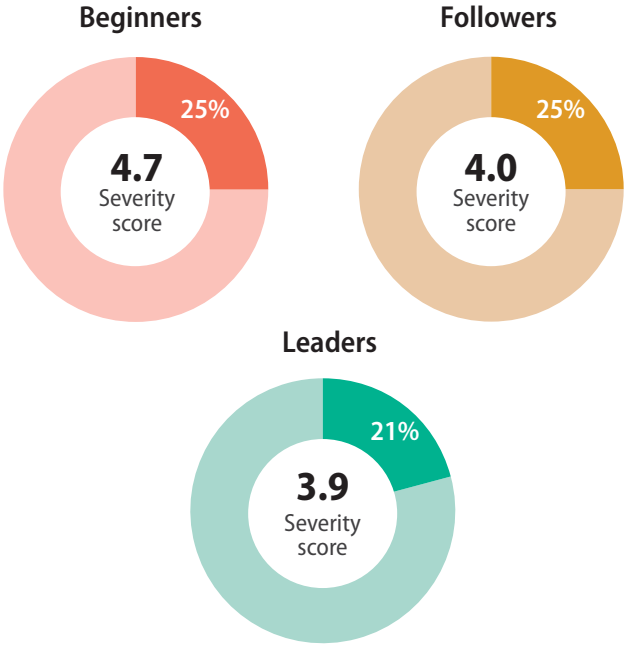

Darker: RAI spending as a percentage of AI spending
Lighter: Remaining AI spend

***N* = 1,502**

Source: Infosys

Unfortunately, hardly any companies are achieving the highest standards set in the RAISE BAR. Indeed, we found that less than 2% of companies can be considered RAI leaders (25 companies out of our 1,500 sample).

A further 15% are in the followers category, and 83%, the vast majority, were beginners that had only met some of the standards. Laggards, those that have implemented very little RAI, represented only nine companies (Figure 10).

It's clear then that many companies have a long way to go to catch up with leaders in RAI, and to gain the benefits of lower severity scores, fewer types of incidents, and lower costs. It's also clear that having experience of a wider range of incidents is linked to improving RAI capability.

This makes it even more important for the vast majority of businesses that are on this journey to learn from those leaders that have gone before — rather than risk learning the hard way themselves.

Figure 10. Only 2% of companies are meeting the best practice RAI standards

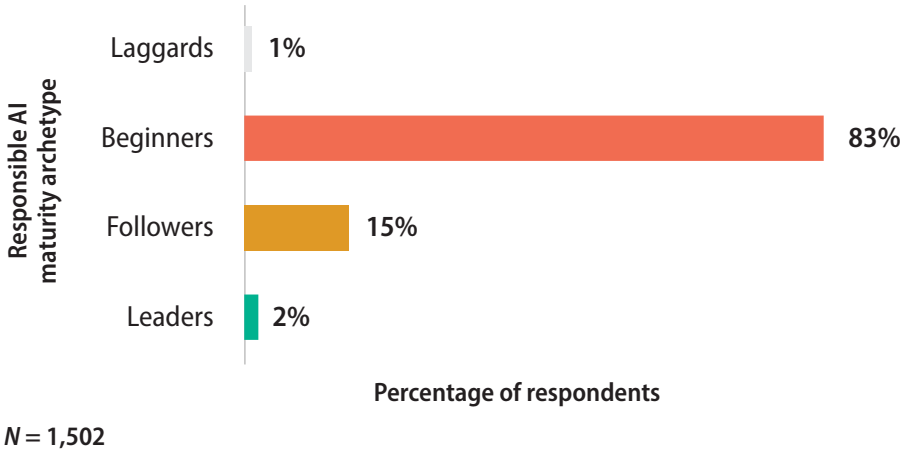

Source: Infosys

Raising the bar

Very few companies are achieving a high level of RAI delivery, but it's imperative that many more eventually reach that level. What can organizations in the follower or beginner categories improve upon to catch up with the leaders? The approach is different for each dimension of our RAISE BAR framework.

Trust

Explainability of AI models is a big element of gaining trust in enterprise AI systems, and therefore affects the impact of enterprise AI products on customers. Studies have shown that 55% of consumers would buy more products from a company whose AI is perceived to be ethical and transparent.

There are several techniques to achieve this, some more powerful than others. For instance, basic techniques like Local Interpretable Model-Agnostic Explanations (LIME) explain single predictions by showing

features that mattered most for a specific result. SHapley Additive exPlanations (SHAP) uses game theory to show how much each feature contributes to a prediction, and is adequate for traditional AI, but in the generative AI and agentic AI era, where models tend to have black-box characteristics, more advanced techniques become crucial.

These techniques include counterfactual analysis (identifying the smallest input changes needed to change a model outcome); chain of-thoughts reasoning (breaking down tasks into intermediate reasoning stages, making the decision process transparent); large language model (LLM)-based evaluation (translating complex model outputs into clear, natural language explanations); and visualization techniques such as principal component analysis (transforming high-dimensional data into 2D/3D plots, revealing hidden patterns and decision boundaries).

According to our experts, companies need to use these more advanced techniques to claim their enterprise AI is explainable. But most do not. While 84% of respondents are using LIME and SHAP, just 43% are doing counterfactual analysis, 39% are doing LLM-based evaluation, and very few (17%) are using visualization techniques.

Similarly, to have reliable enterprise AI systems, organizations need to determine the appropriate level of human oversight, regulatory requirements, technical complexity of enterprise AI systems, and internal governance frameworks and ethical review boards in unison. However, just 5% are doing all three, showing that human oversight is a weak point on the trust dimension.

AI leaders also probe their models with malicious or deceptive inputs to uncover vulnerabilities and assess robustness (known as adversarial testing), and systematically mitigate biases across the entire AI life cycle, something that just 38% of organizations have implemented.

Risk mitigation

Safeguarding against unintended consequences is a key part of enterprise AI risk mitigation. Yet, none of the seven safety measures we asked about in our study — continuous monitoring, anomaly detection, rigorous testing and validation, robust access controls, following ethical guidelines, human oversight, and data quality/integrity measures — have been universally adopted. According to Infosys experts, to be doing RAI well, at least five of these safety measures should be implemented. However, just 4% of organizations have implemented at least five of the seven safety measures, according to our research.

Good RAI, as demonstrated by RAI leaders, also means safeguarding privacy and security. For the security dimension, having a good incident response plan in place is vital, and yet it is only common among our leading companies. Just 48% have grievance and redressal processes in place for when systems malfunction, and roughly two-thirds have a plan to address ethical concerns and isolate

compromised data. Again, all three are foundational for good RAI risk mitigation, and are found at only 10% of organizations.

Data and AI governance

Data and AI governance is more of a bright spot compared with trust and risk mitigation, pushing more companies into the leader/follower categories. A full 83% are validating AI models on unseen data, ensuring that models generalize beyond their training set and reducing the chance of overfit and model bias. Similarly, 89% have documented key parts of the AI life cycle, including datasets, algorithms, and decision-making, which is important when bringing disparate teams together and fixing things before errors accumulate downstream.

Another bright spot is that organizations tend to provide employees with IP awareness training about AI systems (74% are in this category) and have a legal team that advises on AI content, output, and potential infringement (84%). However, RAI leaders are also sticklers for data auditing, checking all third-party data for bias, while continuously monitoring enterprise AI outputs to ensure fairness. Just 33% of organizations were in this category, with a further 65% saying that they did this "often" or "sometimes."

Sustainability

Finally, the sustainability of the whole enterprise needs to be considered, especially as AI's energy demand is projected to surge 160% by 2030. Though a highly contested subject, it is true that most new enterprise AI systems consume significant amounts of energy both in training and inference and are built on GPU-hungry data centers often provided by chip titans like Nvidia. New advancements in engineering have shown significant energy savings in training and inference — demonstrated most effectively by DeepSeek-R1, introduced in early 2025, a Chinese generative model that uses fewer parameters and therefore demanding less energy in computation.

Using these energy-efficient models will be key, as executives build their enterprise AI stack. Keeping environmental impact in mind is important to customers who are increasingly aware of the downsides of fast-expanding AI usage, including its effect on climate change.

In our survey, a full 62% of organizations report monitoring the environmental impact of enterprise AI initiatives. But to be doing well, as demonstrated by enterprise AI leaders, full — not partial — mitigation is necessary, something just 38% of companies are doing. This means developing more compact models that can run on local hardware (instead of rack-mounted hardware), and using data centers optimized on renewable energy sources and advanced cooling systems.

Responsible agentic AI

It's clear from this research that RAI is valued as a critical enabler for enterprise AI, and that companies need to do more to improve the scale, quality, and efficiency of their RAI functions. But there is a further complication — one that could turn all that we know completely on its head.

The emergence of agentic AI not only changes how enterprise AI will be developed and deployed, but it also significantly expands the depth and breadth of use cases that companies could implement. While agentic AI is only emerging today, many believe that it will revolutionize business processes and operating models.

It will also change RAI. Indeed, 86% of executives that are aware of agentic AI believe that it will pose additional risks and compliance challenges to their business (Figure 11).

Infosys believes that to deliver agentic enterprise AI responsibly, a combined product-led, platform-driven operating model needs to be deployed within an organization.

The product-led operating model is central to many digital transformations and includes empowering product teams with the tools and infrastructure to take responsibility for their own outcomes. In such a model, product-centric teams have the ability to build their own applications and processes to best suit the needs of their market and customers — whether internal or external.

Another key element of this model is the use of an objectives and key results (OKR) construct that enables traceability from business objectives right down to product teams. This model works well in the digital era, where centrally commissioned technology solutions are often not agile enough to meet the needs of a business unit.

However, in the world of enterprise AI agents, care needs to be taken before distributing the means to develop and deploy AI solutions across autonomous product-centric teams. Ensuring RAI compliance would be incredibly tough in such a model, where hundreds, if not thousands, of AI agents could be deployed daily.

This is where the platform-driven model comes in. A platform approach involves the creation of a safe environment for enterprise AI agents to be developed and hosted. Enterprise AI agents operate within this environment and have access to preapproved systems and data that are audited to ensure compliance with security, privacy, ethics, and other RAI tenets.

For instance, financial titan ING uses guardrails to preempt regulatory breaches and reputational harm. The company's chatbot embeds RAI to block sensitive data leaks and ensure financial compliance, boosting customer trust. Platforms with embedded RAI also increase employee trust. Microsoft has embedded RAI directly into its development life cycle platform, aligning ethics and trust with product innovation.

As Monika Gupta, partner general manager leading the employee productivity engineering team for Microsoft Digital, says: "When we started out, engineers weren't

Figure 11. Agentic AI will drive innovation but also create new challenges

Percentage of respondents who have heard about agentic AI

***N* = 611**

Source: Infosys

sure what to do with AI and how to do it responsibly, so the default was to restrain their own momentum. Now they know we have RAI built into our practices and our technology as part of an organic process that grows from a root of culture, so they trust the solutions more."

The platform approach also involves the creation of a platform team — a group that has RAI specialists embedded in it, and that works with product teams within the business to build AI agent use cases. They then are supported by an RAI office, which is responsible for setting policies, managing risk across the system, benchmarking processes and agents, ensuring compliance, and researching new risks.

This RAI office, in turn, should be using its own automation tools — which will also involve the use of enterprise AI agents — to manage the large volume of AI use cases that would be built in the platform. They can also use enterprise AI agents to monitor compliance and manage risks.

This product-led, platform-driven structure, outlined in Figure 12 enables innovation and experimentation in a controlled manner, while also embracing the agility of the product-centric operating model.

Figure 12. The agentic AI future: Products, platforms, and centralized RAI governance

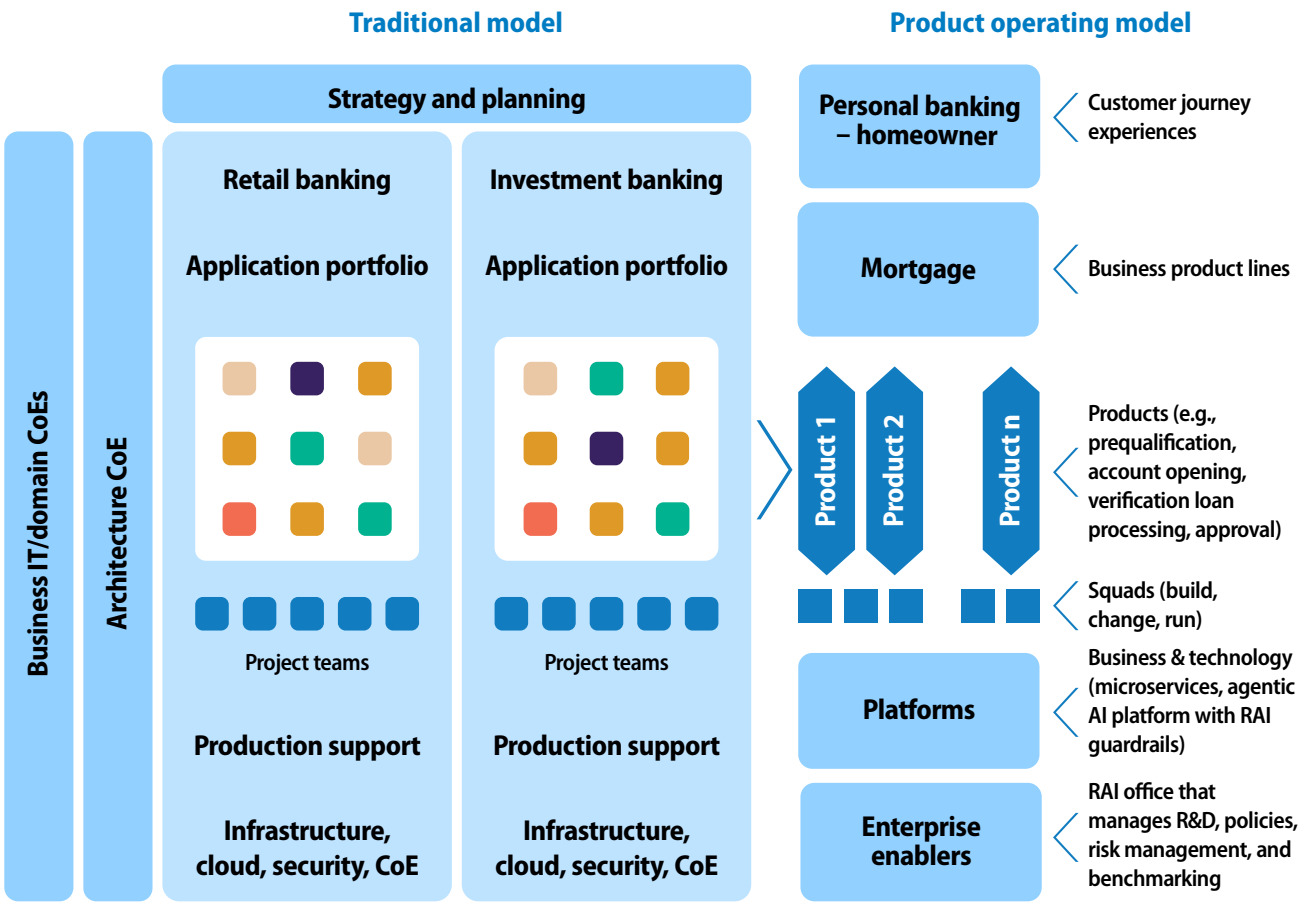

Source: Infosys

From compliance to growth

We recommend that enterprises take these four steps to transform RAI from a compliance function to a growth driver in the agentic AI era:

1. **Learn from the leaders:** Most companies are still only partway on their RAI journey, and to accelerate their capability development, they should learn from organizations that have experienced a wide range of AI incidents. Our data shows that those with higher RAI maturity have encountered more such incidents, which gives them added insight into what techniques they should deploy across trust, risk, data governance, and sustainability to further reduce the cost and severity of RAI.

   Remember, leaders have lower incident costs per initiative, lower severity of damage, and spend less as a proportion of AI budget than trailing organizations.

As organizations learn from leaders, they also increase their reputation in the market for building fail-safe enterprise AI systems, leading to competitive advantage. As Liz Keen, head of clinical governance at Infosys Consulting Australia, says: "Responsible AI confers credibility to companies operating in the health industry."

2. **Combine the product and platform operating models:** Many businesses are already deploying a product-centric operating model, and this is often supported by a platform delivery approach. This shift is also highly conducive to responsible agentic AI development. It meets the needs of product teams and enables better collaboration, control, and risk management. A platform-driven construct also enables shared resources, asset tracking, and ensures enterprise AI progress is documented.

   As noted in Figure 12, the RAI office provides a centralized governance capability that embeds governance into workflows, innovates within guardrails, and ensures that platform teams build compliant infrastructure, such as bias detection application programming interfaces (APIs). By setting companywide policies and risk thresholds, this collaboration prevents silos, as seen in the Infosys hub-and-spoke model, where centralized policy enables decentralized execution. It's also important to embrace modern, automated pipelines to ensure both teams iterate quickly, reduce manual bottlenecks, and maintain high quality in enterprise AI deployments.

3. **Build RAI guardrails in a platform:** To maximize value creation, organizations need to scale RAI techniques without too much added cost. They also need to automate as much as possible, especially with the advent of agentic AI, which, as noted, brings many new risk and compliance challenges. A platform should combine the skills and tools to build AI agents with a library of agents that are made available to enterprise users. This should all be within a secure environment where agents can only access verified data and systems, and guardrails are in place within the design of the agents. It should be supported by a platform team that proactively identifies agent use cases by working with business units and also incorporates RAI specialists within this process. Once these guardrails are in place, enterprises can maximize value creation from their AI initiatives and reduce the 40% of projects within our sample that were either canceled or failed to meet their objectives.

4. **Establish a proactive RAI office:** Effective RAI in the agentic AI era involves continuously evaluating and monitoring emerging vulnerabilities and approaches. Enterprises must regularly upgrade their tools, systems, and processes as AI technology evolves. Infosys' own RAI office has built an automated enterprise AI management system, Infosys AI3S, which includes Scan, Shield, and Steer capabilities.

   Scan helps clients identify the overall risk posture, legal obligations, vulnerabilities, and threats arising due to AI adoption.

The Scan system gathers information from multiple external and internal sources to create a single source of truth for tracking the risk and compliance status of all enterprise AI projects.

Shield offers technical and specialized solutions, guardrails, and accelerators for protecting enterprise AI models from vulnerabilities, and helps embed a responsible-by-design approach across the AI life cycle.

Steer offers advisory and consulting services to help clients become RAI leaders. This means setting up, governing, and managing a dedicated RAI practice, while helping to formulate strategy and expected outcomes via standardized audits and industry certifications.

Using this system, teams can continuously scan for threats and issues, and automatically establish smooth-running enterprise AI operations.

Appendix: Research approach

- ## Survey

  Surveyed 1,500 respondents between March and April 2025 about RAI readiness, implementation and investment, representing businesses with more than $1 billion in annual revenue across seven countries and 14 industries.

- ## Expert analysis

  Interviewed 15 RAI executives to help interpret the data by providing real world examples for trends identified in the survey data.

- ## Model

  We developed a model to measure RAI maturity that we have termed **RAISE BAR** (**R**esponsible **AI S**tandards **E**valuation: **B**enchmark **a**nd **R**eadiness). The model consists of four major dimensions: Trust, risk mitigation, data and AI governance, and sustainability. Within each of these

dimensions, there are sometimes further subcategories. Within trust, we explored explainability, reliability, fairness, and safety. Within risk mitigation, we explored privacy and security. Within data and AI governance, we explored model validation and engineering, intellectual property infringement and protection, and AI audits and standards.

At this level, we developed one to two questions for each subcategory to determine whether a company is performing well or poorly regarding responsible AI in that area. Questions were typically formatted to have an objective answer to help avoid positive bias that can appear in surveys of this nature.

We then created scores for each question based on the answer choices. Question scoring can be grouped into three distinct categories:

- The first category requires respondents to have selected a core group of practices as well as another practice beyond the core to be considered "doing well."

- The second category requires a minimum number of practices (e.g., five of seven) to be considered "doing well."

- The third category is a simple gradation of a single selection question, where only the top one or two answer choices qualify for "doing well."

We then created thresholds for points based on the "doing well," "doing core components," "missing some core components," and "minimal effort" framework used in scoring the individual questions.

- ## Insight

  We identified that very few companies meet most of the standards laid out for responsible AI, with less than 2% meeting all the standards tested by RAISE BAR, and only a further 15% being most of the way toward becoming a leader.

- ## Outcome

  Identified a correlation between RAI maturity (as defined by RAISE BAR scores) and the following:

  - AI incident damage severity
  - Number of AI initiatives
  - Number of types of AI incidents
  - Spending on RAI as a proportion of AI implementation spending

  These correlations can be explained in several ways, as we have detailed in the report.

Authors

Samad Masood | Infosys Knowledge Institute, London
Harry Keir Hughes | Infosys Knowledge Institute, London

Contributors

Syed Ahmed | Head of Responsible AI Office, Infosys
Mandanna Appanderanda | Head of Responsible AI Office, Americas, Infosys
Gaurav Bhandari | Associate vice president, AI and data strategy, Infosys
Gary Bhattacharjee | Vice president, AI and data strategy, Infosys
Ritarshi Chakraborty | AI Center of Excellence, Infosys
Rajeshwari Ganesan | Distinguished technologist, machine learning and AI, Infosys
Liz Keen | Head of clinical governance, Infosys Consulting
Mukul Khare | Delivery manager, Infosys
Dr. Catherine Lopes | AI and data, strategy and advisory, APAC, Infosys
Shalini A. Nair | Vice president, corporate initiatives, Infosys
Rahul Pareek | Principal consultant, responsible AI, Infosys
Kaushal Rathi | Senior associate analyst, AI, Infosys
Saibal Samaddar | AIX lead India, Infosys Consulting
Jayprakash Singh | Client services, Infosys
Srinivasan Sivasubramanian | Industry principal, Responsible AI Office, Infosys
Alok Uniyal | SVP and head, enterprise quality solutions, Infosys

Analysis and production

Isaac LaBauve | Infosys Knowledge Institute, Dallas
Pramath Kant | Infosys Knowledge Institute, Bengaluru
Pranav Tekade | Infosys Knowledge Institute, Bengaluru
Sandeep | Infosys Knowledge Institute, Bengaluru
Kate Bevan | Infosys Knowledge Institute, London
Pragya Rai | Infosys Knowledge Institute, Bengaluru

About Infosys Knowledge Institute

The Infosys Knowledge Institute helps industry leaders develop a deeper understanding of business and technology trends through compelling thought leadership. Our researchers and subject matter experts provide a fact base that aids decision making on critical business and technology issues.

To view our research, visit Infosys Knowledge Institute at infosys.com/IKI or email us at iki@infosys.com.

For more information, contact askus@infosys.com

Infosys®
Navigate your next