CASE STUDY

A STRAW TO BRICK GOVERNANCE STORY

How an International Financial Enterprise pivoted to leadership in Responsible AI

Abstract

Recognising the urgent need for a unified framework for Responsible AI, Aria Portman, Senior Manager for AI Governance at an international financial enterprise, partnered with Infosys BPM. Infosys BPM together with Infosys Responsible AI Office, deployed technical guardrails and an automated monitoring system, which strengthened risk and compliance postures, gave rise to ethical and trustworthy AI systems, and positioned the enterprise as a leader in Responsible AI.

Infosys®
Navigate your next

When governance gets suddenly challenging

Aria Portman had been working tirelessly on a proposal for days, one that would allow her to usher in a new era of governance in artificial intelligence (AI) at the international financial enterprise where she served. As Senior Manager for AI Governance, she headed their Responsible AI practices, which included managing technical guardrails, ensuring the ethical deployment of AI systems, and protecting the enterprise against vulnerabilities.

The financial enterprise desperately needed a unified framework for Responsible AI, a need that became more evident with each AI implementation. And so, Aria hoped to draw the attention of the top management to this matter through a persuasive proposal. She knew that although AI could solve complex problems with ease, it could do more harm than good without governance and risk management frameworks.

Her premise was straightforward: the absence of technical guardrails and ethical oversight had left the enterprise's AI systems vulnerable to security, compliance, and user trust-related risks. The machine learning models were prime targets for data poisoning, prompt injection, adversarial attacks, and other advanced threats. The danger was that these exposures could compromise the integrity of AI outputs. Unreliable business decisions always led to financial losses.

Further, the risk of AI models generating harmful or biased content was high as the enterprise lacked sufficient controls to govern bias, transparency, and accountability. Aria emphasized that, since the enterprise had significant influence in global markets, the limited transparency and accountability made it challenging to justify AI-driven outcomes in regulated industries. The lack of auditability also opened the door to operational inefficiencies and legal exposure while affecting user confidence and reputation.

Aria also explained the connection between regulatory compliance and technical safeguards. She stated that the absence of guardrails made it harder to detect and prevent malicious activities. In the event of any AI threat, such as a breach in privacy, the enterprise would be fined for non-compliance with data protection laws, eroding stakeholder trust once again.

As the internal teams lacked the necessary experience, she proposed seeking expert assistance to create an AI governance framework. The enterprise was already struggling with inconsistencies in risk management, a slowdown in innovation, and an increase in vulnerability to emerging AI threats. Furthermore, a general reluctance to change stemming from concerns about complexity and disruptions delayed the adoption of critical AI controls and increased their risk.

The management team reviewed the proposal and extended their approval to engage an expert at the earliest. So, Aria started searching for a suitable partner who could help her design and implement the necessary AI guardrails. The search led her to Infosys BPM, whose experience and expertise in Responsible AI were in line with her requirements. After consulting with internal stakeholders, Aria negotiated a formal partnership with Infosys BPM.

Rebuilding before danger strikes

Arjun Das, an experienced program manager from Infosys BPM along with Infosys Responsible AI Office, quickly assembled a cross-functional team whose experience spanned Responsible AI, data security, compliance, and AI ethics. The team started with an in-depth analysis of all internal AI applications, mapping business processes, data flows, and potential vulnerabilities. Using the insights from their analysis, they designed a tailored Responsible AI framework. The team also suggested certain technical guardrails that aligned with the enterprise's use cases and risk profile.

Approach summary

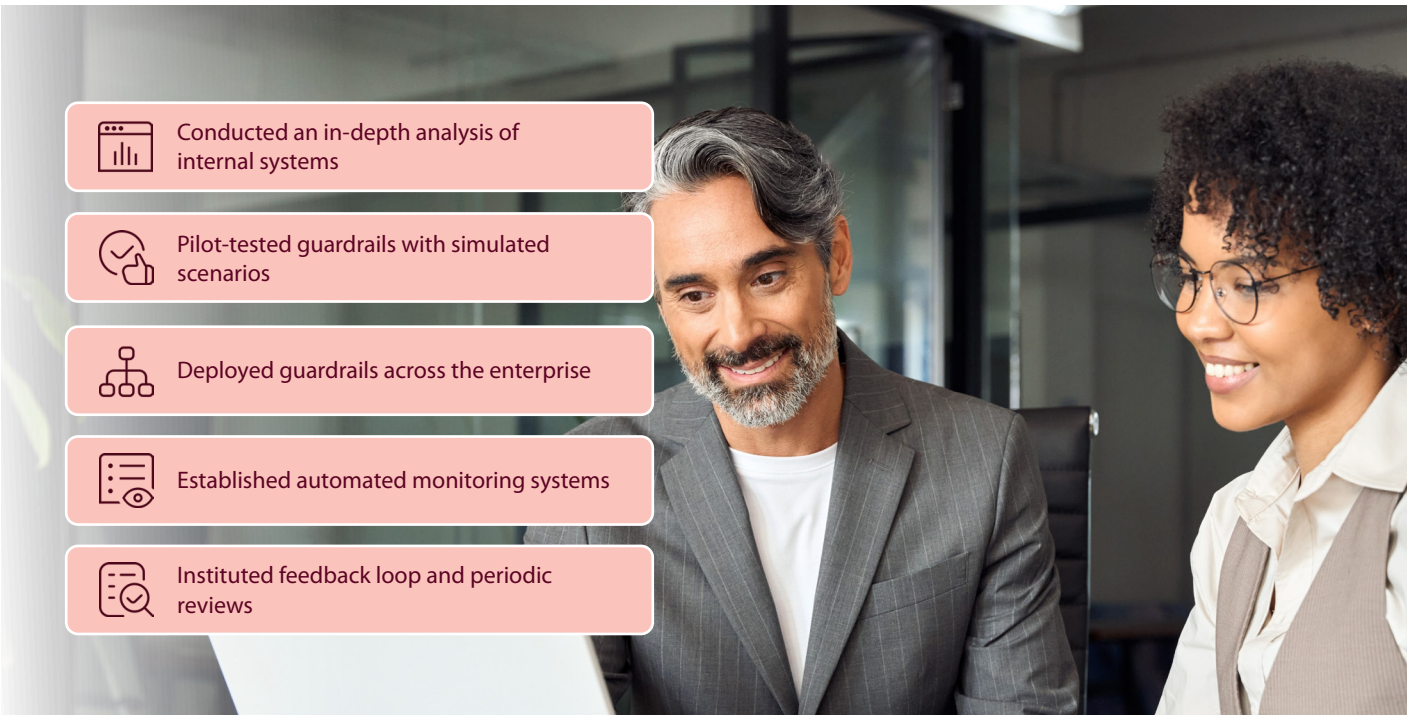

The recommendation concerned a comprehensive AI-driven suite of services, solutions, and platforms powered by generative AI. The suite featured more than 150 pre-trained models and over ten platforms. It prioritised ethics, trust, privacy, security, and compliance and was anchored in a responsible design framework.

However, some stakeholders raised a few concerns over the potential for disruptions to workflows and the complexity of implementing advanced guardrails. Aria was also keen to know how the Infosys BPM team would handle issues with data integration and manage accountability for AI-based decisions. In response, Arjun worked closely with the stakeholders to resolve their concerns and put a phased deployment plan into action.

A pilot was deployed in a controlled environment and tested using simulated adversarial scenarios to validate the effectiveness of the guardrails. Once the pilot delivered the desired results, the team rolled out the framework across all relevant AI systems and business units, fully integrating it into the existing IT infrastructure.

While Infosys BPM was initially engaged only to design and deploy technical guardrails, the scope was later expanded. In line with the new requirements, Arjun oversaw the establishment of automated monitoring systems that provided ongoing risk assessments and incident detection capabilities. Workshops and training sessions that followed helped build awareness and prompted buy-in from stakeholders. Arjun also created a feedback loop for continuous improvement and instituted periodic reviews. This step was crucial in helping the enterprise stay ahead of technological and regulatory developments.

Throughout the implementation, the Infosys BPM team maintained clear documentation and engaged with stakeholders. Furthermore, the phased execution plan helped mitigate delays and close gaps in coverage, allowing Arjun to continuously fine-tune the guardrails as they moved ahead.

A structure that holds fast

Aria met for a detailed review with Arjun when his team completed the final phase. Starting with the technical guardrails, Arjun highlighted how the implementation had strengthened the organization's AI risk posture and significantly reduced any exposure to adversarial threats. With the new risk mitigation system, the enterprise could now identify and reduce bias, as well as prevent privacy breaches and other security threats. The guardrails also protected against copyright infringement, reducing the risk of legal hassles.

Key benefits

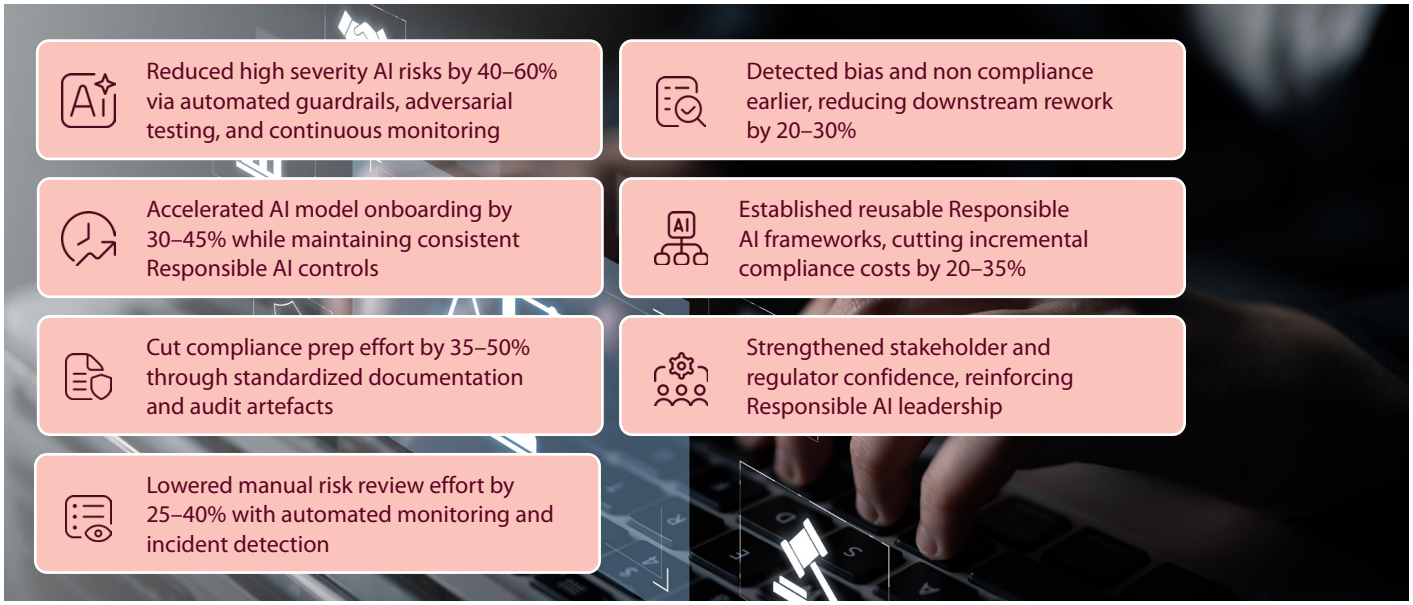

An added moderation layer helped protect against malicious use by flagging harmful or biased AI outputs before they reached users. Arjun also pointed out that the automated monitoring system had streamlined risk management processes, which had a direct impact on operational efficiency. Together, the proactive risk mitigation measures and the reusable framework established a foundation for the growth of Responsible AI, reducing future compliance costs.

Aria noted that as compliance with industry standards improved, it strengthened the enterprise's reputation and boosted stakeholder confidence. In addition to a sustainable AI innovation process, the enterprise also gained ethical and trustworthy AI systems embedded with the principles of fairness, inclusivity, and accountability, all of which would go a long way towards fostering user trust.

As the review concluded, Aria and other stakeholders acknowledged that the skill and expertise of Arjun, Infosys Responsible AI Office and Infosys BPM team were crucial in implementing such an initiative. After all, they had earned the enterprise the prestige of being a leader in responsible AI practices. A new era had truly begun at the enterprise, and Aria knew she had a trustworthy partner to rely on as she looked towards the future.

**Names have been altered to preserve the identities of the people involved.*

For more information, contact infosysbpm@infosys.com

Infosys®
Navigate your next