Standard Bank
Corporate & Investment Banking

SBSA

RESPONSIBLE AI APPROACH

RESPONDING RESPONSIBLY TO THE COMPLEXITY AND ACCELERATING IMPACT OF ARTIFICIAL INTELLIGENCE

STANDARD BANK'S APPROACH

RESPONDING RESPONSIBLY TO THE COMPLEXITY AND ACCELERATING IMPACT OF ARTIFICIAL INTELLIGENCE – **STANDARD BANK'S APPROACH**

Generative AI is advancing at a pace that is fundamentally reshaping how decisions are made, services are delivered, and value is created. This rapid acceleration brings significant opportunity, but also increased uncertainty and risk. As AI systems become more complex, interconnected, and less transparent, the ability to fully understand, control, and predict their behaviour continues to diminish.

In this environment, **responsibility can not be assumed. It must be deliberately designed, actively governed, and continuously revisited.** Responsible AI is therefore not about constraining innovation, but about sustaining it—ensuring that technological progress remains aligned with human values, institutional trust, and long term societal benefit. Achieving this requires purposeful partnerships across academia, industry, government, and society, because responsible innovation is strongest when it is shaped together.

The strategic partnership between Standard Bank and the School for Data Science and Computational Thinking at Stellenbosch University reflects a forward-looking collaboration between industry and academia. By combining Standard Bank's financial and industry expertise with the School's strength in data science, computational thinking, and AI, the partnership advances innovative research, practical solutions, and skills development for the next generation of data-driven leaders. To advance responsible AI, a specialised working group was convened, bringing together the Standard Bank Lab at the School for Data Science and Computational Thinking, the Department of Philosophy, and the Standard Bank CIB team. Over several weeks, the group engaged collaboratively to explore and shape practical approaches to responsible AI.

One of the assets developed by this working group is the ethical AI framework to guide responsible innovation. The framework establishes a clear foundation for designing and deploying AI systems that are safe, robust, and secure, with appropriate human oversight and accountability throughout their lifecycle. It emphasises transparency, explainability, and interpretability, while preserving privacy, promoting fairness, and actively preventing discrimination.

In Standard Bank, these principles are applied consistently across selected high-impact use cases of AI that support our clients, operations, and risk management activities. AI is used to enhance client engagement and service delivery through digital assistants and client support tools, as well as to enable more personalised interactions and

onboarding experiences. Within risk, finance, and compliance functions, AI supports activities such as transaction and fraud monitoring and early identification of potential risk indicators. We also apply AI to improve operational efficiency through intelligent document processing, automated classification, and the synthesis of reports and insights, while supporting relationship managers with tools that surface relevant client insights and AI augmented banking. Capabilities such as SmartNudge, our inhouse built AI recommendation engine for clients, further demonstrates how these AI-enabled insights and experiences can be delivered directly to our clients, while maintaining transparency and responsible data and AI usage.

The framework is informed by industry standards, policy models, governance proposals, and global best practices, and has been purposefully tailored to address the complexities of innovating within a broader ecosystem of clients, partners and stakeholders.

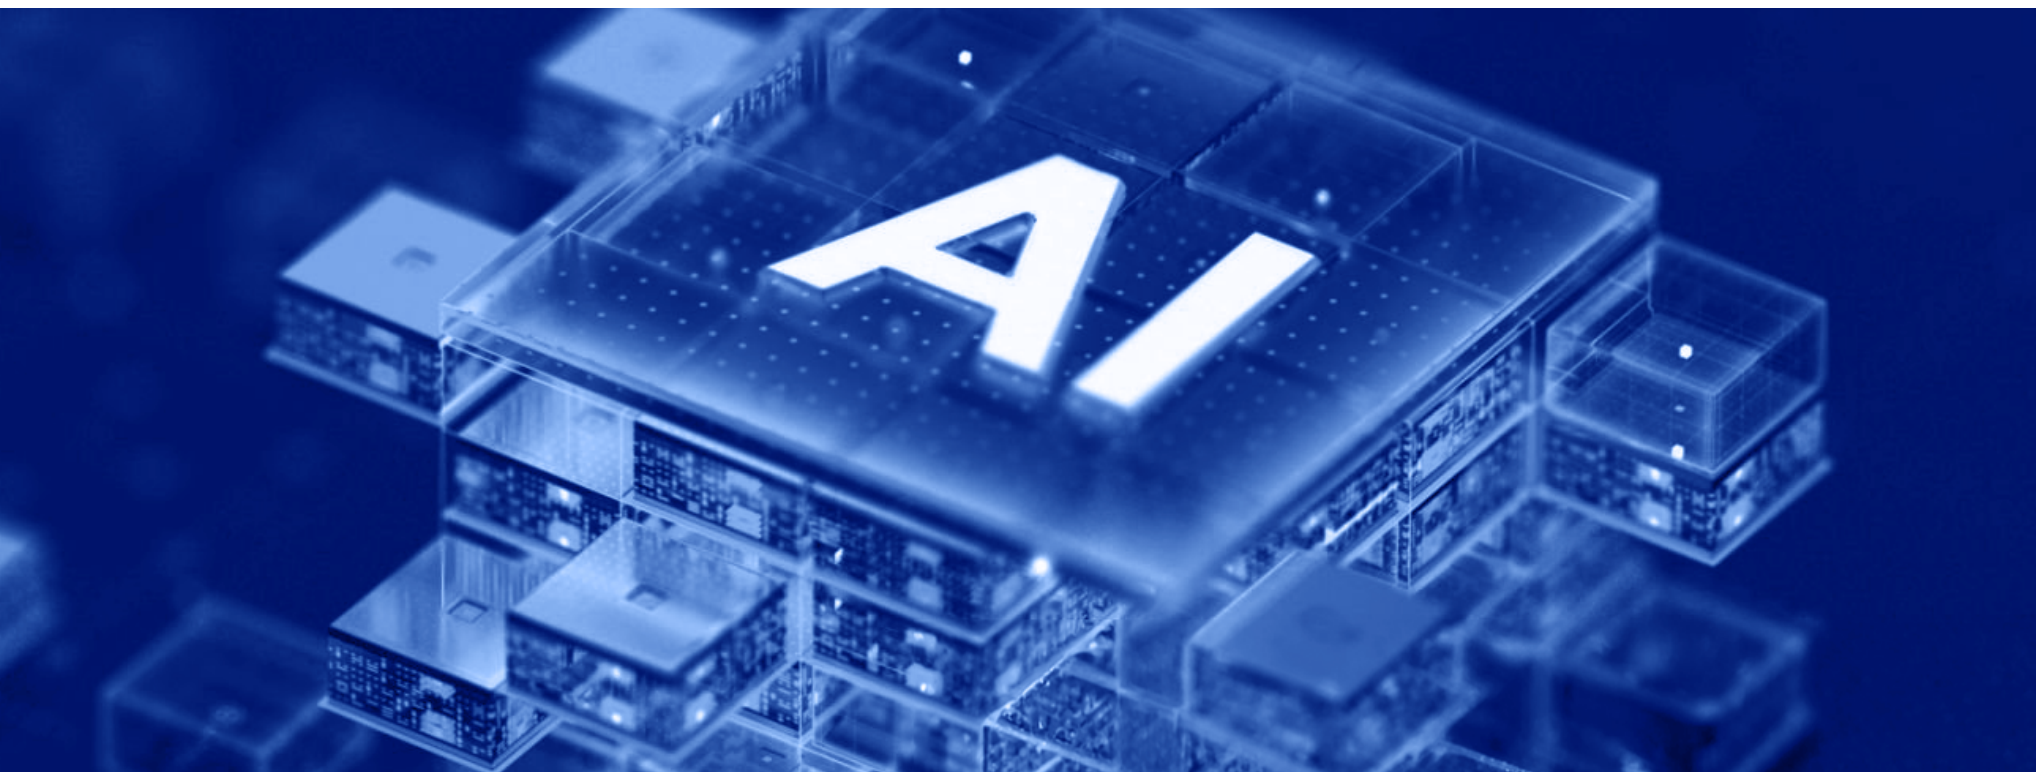

Operationalisation of Responsible AI framework

To operationalise our ethical principles, we have developed questions appropriate to each phase of the product development lifecycle. The questions are intended to prompt thoughtful deliberation, encourage documentation of reasoning, and provide a record of the ethical considerations that informed decisions across the product lifecycle.

The question sets are designed to be applied as a continuous, iterative cycle across the full product lifecycle, from early idea formation through to deployment and ongoing use. The phases—discovery and idea validation, development, and application—are deliberately separated to reflect how ethical considerations evolve as a product moves from concept to

reality, while also reinforcing that ethics is not a one off checklist, but an ongoing discipline that requires revisiting decisions as new information, risks, and impacts emerge.

The **Discovery and Idea Validation** phase focuses on whether a product should be pursued at all. The considerations in this phase are intended to surface foundational ethical risks and value judgements at the earliest possible stage, when changes are least costly and most effective. They prompt teams to examine necessity, alignment, benefit, harm, fairness, bias, consent, and stakeholder impact before significant resources are committed. In this phase, the emphasis is on ethical justification, problem framing, and informed decision making, ensuring that the product concept aligns not only with business objectives but also with human, societal, and environmental values.

For instance, in customer relationship management, GenAI-driven customer insights can support more personalised, transparent, and responsible engagement when guided by strong privacy, consent, and fairness principles. Additionally, personalised client engagement and marketing can strengthen relevance, trust, and customer experience when designed to be fair, inclusive, transparent, and respectful of customer choice.

In the validation stage, GenAI tools are employed to convert ideas to prototype solutions as well as synthesise initial data points in order to achieve a fair outcome. In addition, validation ensures that the proposed solution is feasible in solving the problem.

The Discovery phase should guide whether a product moves forward, is refined, paused, or rejected by carefully weighing its ethical value, stakeholder impact, and alignment with responsible AI principles. At its core, this phase seeks the right balance between potential harm, bias, and the broader benefits the project aims to deliver.

The **Development** phase addresses how the product is designed, built, and governed once a decision to proceed has been made. The considerations here translate high level ethical principles into concrete practices across data use, model design, system architecture, human oversight, privacy, security, fairness, transparency, accountability and auditability. This phase recognises that many ethical risks—such as bias, opacity, loss of autonomy, or unsafe behaviour—are introduced through technical and institutional choices. Accordingly, the questions are structured to help teams embed ethics into day to day development decisions, governance processes, and controls, and to ensure that products remain explainable, auditable, and responsive to risk throughout their lifecycle. These have been embedded in practice through careful model design, governed data pipelines, and control frameworks in client-facing solutions such as SmartNudge.

Another illustrative example, in the FMCG sector, an AI-based solution for stock and inventory management needs to be properly designed with the user who is going to own and utilise the final solution fully part of the process. The design must rely on data that is reliable, and

representative of all relevant operational factors whilst ensuring proper access and controls are in place. This supports technical robustness and safety by ensuring that the system produces dependable and fair outputs. Assessing the robustness of such a tool requires careful consideration of the model design, quality of data inputs, and access management controls. For instance, if the goal is to solve for stock forecasting, the data will need to be modified to be AI-ready to ensure easier integration into an AI-based forecasting tool.

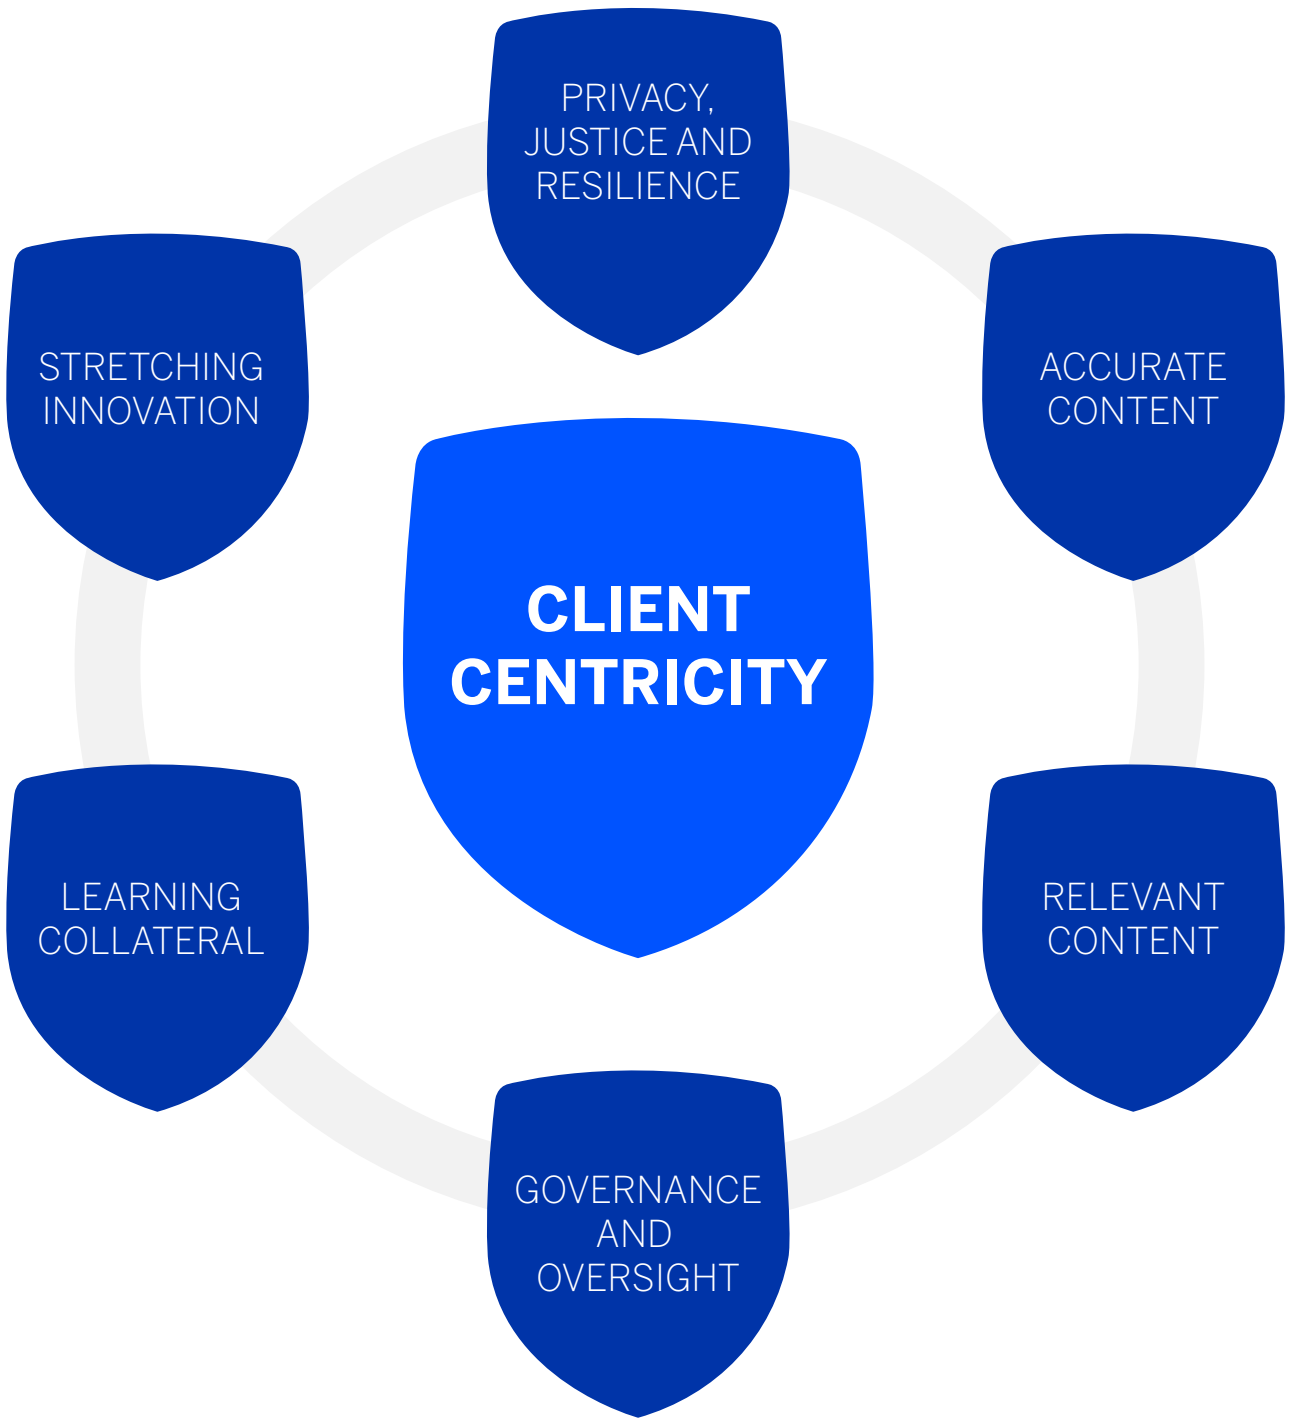

Figure 1: Evaluation checklist post AI-Solution development.

The **Application** phase focuses on how the product is brought to market, used, supported, and maintained over time. The considerations in this section recognise that ethical responsibility does not end at launch. Instead, they emphasise honest communication, expectation management, accessibility, fair support, performance monitoring, feedback, remediation, and continuous improvement. This phase ensures that real world use aligns with what was promised during discovery and development, that users retain meaningful autonomy and access to human support, and that emerging harms, drift, or inequities are detected and addressed.

For instance, after development, the final AI-based solution could be a tool, a workflow redesign, an independent solution, or an AI agent that is surfaced to the final user for consumption. To this end, considerations such as user onboarding to the solution, support training materials, operational or business risk classification of the solution, explainability of the output, and ongoing support for the user needs to be clear and well defined.

One may further consider other AI tools to help with monitoring the AI-based solutions such as using another AI-as-a-judge or other AI tools to continually observe and monitor degradation of the AI-based solution.

Figure 2: Key considerations and recommendations.

Taken together, these phases and questions provide a structured way of thinking about responsible AI, but also illustrate that such a framework cannot be treated as a static governance document. It must be understood as a living framework that evolves alongside the innovations it governs. As the pace of AI innovation increasingly exceeds the ability of traditional policies, regulations, and business processes to adapt, organisations face the ongoing challenge of remaining both innovative and responsible. Meeting this challenge requires continuous review, learning, and adaptation, as well as a willingness to question established assumptions and operating models. **By approaching responsible AI as an enduring discipline rather than a fixed destination, organisations can better navigate uncertainty, sustain trust, and ensure that innovation remains aligned with ethical, societal, and long term value considerations.**

Disclaimer:

Unless expressly stated otherwise, all content contained in this report is protected by copyright and is owned by SBSA. Although SBSA has taken reasonable care in preparing this report, it makes no representation or warranty, express or implied, as to the accuracy, completeness or currency of the information contained in it. This report is provided for general informational purposes only and does not constitute, and should not be construed as, legal, investment or other professional advice. It has not been prepared with regard to any person's individual circumstances and does not constitute a personal recommendation. To the fullest extent permitted by law, SBSA accepts no liability for any direct, indirect, consequential or other loss or damage arising from or in connection with the use of, or reliance on, this report or any information contained in it. Recipients must independently verify the information and obtain advice from appropriately qualified experts or professionals before making any decision based on it. The report must not be relied on as a substitute for the exercise of independent judgment.