"A DECISION-MAKING FRAMEWORK FOR PRIORITIZING AI ADOPTION ACROSS ENTERPRISES”

Research Paper

dr. Gustav Lindéus, SSBM, Geneva, Switzerland, gustav@ssbm.ch

dr. Deepak Kota, SSBM, Geneva, Switzerland, deepakk@ssbm.ch

“Abstract”

This paper addresses the increasing challenge of how to select and prioritize generative AI initiatives when technologies develop faster than governance, compliance, and funding processes. We propose the GAIQ framework: a design-science-based, gate-driven model for qualifying GenAI use cases along three dimensions, namely, PVI, TFR, and ERC. The model structures decision-making through a SEA sequence of scanning, evaluation, and activation and applies weighted thresholds so that use cases that are strategically attractive but weak in ethics or technology cannot advance. Two complementary instruments, NEXA and NOVA, extend the framework to investment decisions. Validation on simulated enterprise scenarios shows that GAIQ produces more consistent, auditable, and business-aligned recommendations than generic AI maturity models, thereby closing the gap between high-level AI strategy and operational implementation.

Keywords: GenAI, AI, GenAI governance, use-case prioritization, gate-based evaluation, risk-adjusted ROI

1 Introduction

Generative AI (GenAI) is a subset of artificial intelligence (AI) that focuses on creating new, unique content from existing data (IBM, 2024). It uses algorithms and models to generate text, images, music, code, and other media that are original yet reflect the patterns and characteristics of the input data (Amazon Web Services, 2025.; McKinsey Global Institute, 2023). AI has become a key driver of economic growth and competitiveness, yet its transformative potential depends heavily on governance quality as well as technological infrastructure (Lindéus and Shetty, 2024). Recent regulatory debates emphasize balancing opportunity with governance. The EU’s AI Act follows a risk-based approach but still lacks a proper risk–benefit analysis and solid empirical support (Ebers, 2024). The selection, evaluation, and prioritization of GenAI use cases present significant challenges for organizations due to inconsistent decision-making (McKinsey & Company, 2023), resource allocation difficulties (Deloitte AI Institute, 2024), and a rapidly evolving technology landscape (Ernst & Young, 2024a).

Current methods are rooted in subjective criteria, and this leads to arbitrary choices that have no alignment with organizational objectives (Ransbotham et al., 2019). Without abundant resources and abilities, good prioritization is crucial; otherwise, organizations risk expending valuable time and financial inputs on the wrong projects. The rapid pace of AI technology development complicates it even more, and teams struggle to evaluate new instruments and methods effectively (Salesforce, 2023). This lack of systemic decision-making does not merely generate mediocre outcomes and increased risk exposures but also hinders agility, rendering it impossible for organizations to react fast to market changes (Baxter and Schlesinger, 2023). To address such issues, a systematic approach that establishes crisp selection criteria, encompassing risk management, and mapping selected use cases to strategic goals is not only beneficial but a sheer necessity for organizations. It will enhance decision-making in an informed way and stimulate innovation in GenAI initiatives.

The paper will help organizations avoid pitfalls in GenAI implementation by providing an organized procedure for the identification and prioritization of GenAI application instances. It emphasizes the need to ensure that GenAI projects align with the overall company objectives and determining the ways in which GenAI can help drive them. Companies are faced with serious challenges in defining, examining, and ranking business application instances for GenAI implementation. Challenges include insufficient technical expertise, data privacy concerns, integration issues with current systems, and the expense of AI solutions. Consequently, companies cannot use GenAI adequately, which results in wasted data and missed opportunities for innovation.

To address these issues, the introduction of the new Integrated Value Evaluation Model for Gen AI Use Cases, “GAIQ” is crucial. This tool has a gated process that critically tests use cases to see that they meet all the necessary considerations, e.g., technical feasibility, financial viability, protection of data, ethics and business alignment, before embarking on implementation. This methodical process helps companies make sound as well as informed decisions, mitigate risk, and obtain the maximum capabilities from GenAI technologies.

Recent industry feedback confirms the urgency of adopting structured frameworks. Many organizations are already utilizing GenAI to some extent for processes such as document creation, customer interaction, and one-to-one marketing. But they are stuck on areas such as the identification of ROI, integration, and setting priorities. Informal methods like RICE scoring or internal triages offer very limited consistency and effectiveness, so there is clearly a need for a more standardized approach.

By focusing on those gaps, the GAIQ framework helps organizations to navigate through the complexities resulting from GenAI decision-making and promote innovation while prudently managing risk and resources. Its benefits over traditional approaches have been attested to by industry experts, which again underlines its importance and impact.

2 Literature Review

This literature review critically reviews the current knowledge of GenAI adoption in business strategy and implementation. It synthesizes theoretical foundations, practical challenges, and the building of decision frameworks, with focus on the evaluation and prioritization of GenAI applications. While GenAI has progressed significantly, the literature reveals a core deficiency: the absence of properly structured, question-based systems that help organizations select and implement high-impact AI projects.

2.1 Theoretical underpinnings and business impacts

The relevance of GenAI to business strategy stems from understanding its theoretical foundations, according to recent research (Gupta, 2024). GenAI is rapidly becoming increasingly popular across numerous sectors on the premise of maximizing efficiency, customer interaction, and return on investment (ROI). GenAI significantly maximizes ROI and allows brands to engage more effectively with consumers (Patil, Rane and Rane, 2024a). But difficult tasks such as integration complexity and ethical concerns have to be addressed in order to harness its complete potential.

2.2 Challenges and ethical concerns

There is massive potential for GenAI to transform sectors with the automation of tasks, improving creativity, and fostering innovation. Its use is not without risks, however, including ethical concerns, data privacy, and the ability to generate erroneous or biased content (Baxter and Schlesinger, 2023). Applications prioritization assists in mitigating these risks and realizing effective implementation. In spite of pervasive AI investments, 70% of companies state no or minimal impact, and even where there are significant investments in AI, 40% still find no business benefits (Ransbotham et al., 2019).

2.3 Use case prioritization

Selecting the right use cases for GenAI is critical to ensure maximum business value. Van der Veen (2024) provides an organization readiness framework to deploy GenAI, with a focus on use case prioritization and end-to-end business impact assessment. Patil, Rane, and Rane (2024b) continue to discuss the importance of assigning high priority to effective use cases in order to realize significant business outcomes and actual ROI approximations.

2.4 Technical, ethical and cybersecurity challenges

Implementation of GenAI involves addressing technical, ethical, and cybersecurity challenges. Patil, Rane, and Rane (2024c) refer to the necessity of investing in robust cybersecurity controls to mitigate risks associated with GenAI, which will influence ROI calculations. Wala and Wooten (2024) highlight that the key challenge to large enterprise is not the technology but how effectively to prioritize use cases and adequately conduct business impact analysis to ensure ROI.

2.5 ROI calculations

It is challenging to precisely calculate ROI on GenAI initiatives because companies must forecast future implications and justify outlays. Sterne (2023) discusses a set of methods for calculating ROI and the constraints of forecasting business value for GenAI, underlining concentrating on use cases with unquestionable business implications. Rajaram and Tinguely (2024) provide a practical guide to help SMEs break GenAI adoption obstacles and provide recommendations that can be employed to facilitate prioritization initiatives, to the ROI.

Zao-Sanders (2025) provides an updated overview of the application of GenAI in personal and commercial environments. The study highlights the emergence of Custom GPTs tailored for specific requirements, new competitors like DeepSeek and Grok, and innovations such as Google's podcast generator, NotebookLM. The article emphasizes the broadening access to GenAI and the reduction in costs, which have significantly impacted its adoption and utility. Global Lenovo research reveals that proving ROI remains the greatest obstacle to AI adoption despite higher spending (Lenovo, 2025).

2.6 Practical recommendations

Master et al. (2024) provide realistic recommendations on businesses leveraging GenAI. They emphasize the necessity of thorough business impact analysis and accurate ROI calculation, as needed to prove value of investments in GenAI. Their paper outlines how companies can apply GenAI offensively and cope with the implementation challenges.

2.7 Ethical guidelines and innovative business models

Kalusivalingam et al. (2022) discuss the application of generative adversarial networks (GANs) and reinforcement learning (RL) in formulating new business models. They observe that ethical principles have to be adhered to and comprehensive impact analyses need to be conducted to facilitate the effective use of these technologies. Their work underscores the revolutionary potential of GANs and RL to infuse strategic transformation in firms, provided that ethical issues are adequately addressed.

2.8 Relevant frameworks

The deployment of GenAI use cases represents a major strategic and operational challenge: not all use cases provide equal value, and most companies do not know how to identify and prioritize those that will have the greatest effect. If not effectively evaluated, companies can squander resources on initiatives that do not contribute to overall strategic goals. Gartner (2023) reports that 77% of CEOs believe AI is ushering in a new era of business change, but many believe that their technology leaders are not ready to turn these changes into business results. At the same time, 82% of technology leaders polled by EY plan to expand their AI investments in the coming year, highlighting the increasing imperative of successful AI planning and deployment (Ernst & Young, 2024b). Structured frameworks can aid organizations in overcoming these difficulties by offering a guide for assessing, choosing, and scaling GenAI use cases. These frameworks are decision-support tools that map technological capabilities with business value and risk dimensions.

Framework	Strengths	Gaps / Limitations
PwC's GenAI Value-Realization Flywheel (Greenstein, Light and Likens, 2024)	Structured approach; focus on value creation; iterative learning process	Requires significant initial investment; lacks practical implementation guidance
McKinsey's AI Transformation Framework (McKinsey & Company, 2023)	Holistic approach; aligns with business strategy; promotes cross-functional collaboration	Lacks detailed, actionable steps; may overwhelm new AI adopters
Deloitte's AI Framework (Deloitte, n.d.)	Thorough structure; includes feasibility and impact assessments; promotes responsible AI	Needs significant customization; limited step-by-step implementation guidance
Gartner's AI Maturity Model (Gartner, 2023)	Clear maturity levels; comprehensive evaluation metrics; strong emphasis on risk management	High-level abstraction; resource-intensive to achieve higher levels of maturity

Table 1. Summary of Prominent GenAI Frameworks.

PwC's GenAI Value-Realization Flywheel adopts an iterative approach centred on value creation, helping organizations move from ideation to scaling (Greenstein, Light and Likens, 2024). It focuses on continual learning and systematic advancement but requires high up-front investment and is not highly specific about operational execution.

McKinsey's AI Transformation Framework provides a high-level view that integrates AI strategy with business objectives, stressing cross-functional alignment (McKinsey & Company, 2023). Its strategic abstractions, however, can be counterproductive to adoption, especially for companies developing their AI capabilities, and it provides no specific guidance on how to prioritize certain GenAI projects. Deloitte's AI Framework integrates feasibility and impact assessments and places heavy importance on ethical AI practices (Deloitte, n.d.). It offers an overall framework but is often argued to necessitate heavy customization and the lack of step-by-step implementation guidelines, especially for cost-constrained teams. Gartner's AI Maturity Model enables organizations to gauge advancement through levels of maturity and concentrates on risk governance and quantification (Gartner, 2023). However, its focus is still diagnostic and not directive, and progress towards greater maturity tends to be resource intensive. The illustration below provides a quantifiable comparison of the different frameworks, highlighting e.g. the gap in the operational guidance.

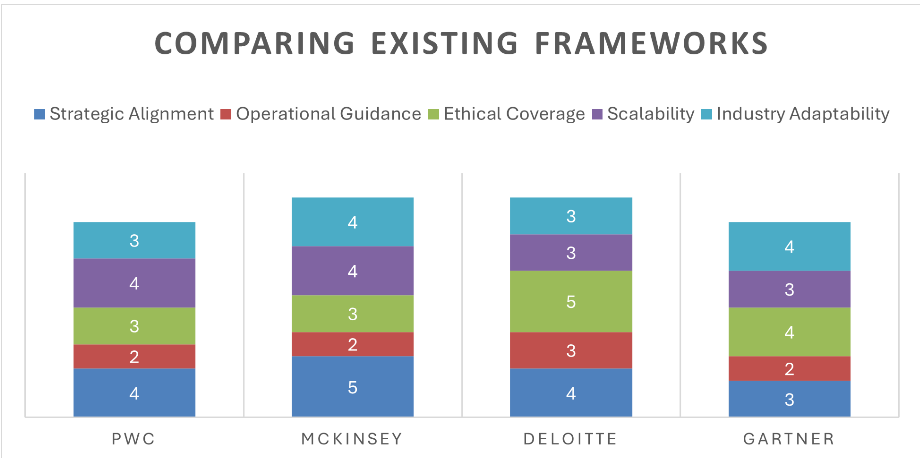

Figure 1. Comparison of existing frameworks, based on a scale from 1-5, where 5 is the highest, reflecting their strengths and weaknesses.

Recent academic literature confirms these concerns. Nguyen (2025) studies as well as explores about the ethical and pedagogic constraints of GenAI in education and demands clear guidelines and constructing AI literacy. Similarly, Samala et al. (2024) propose a taxonomy of GenAI use cases and underline the importance of ethical guardrails and informed evaluation practices. Across both academic and practitioner literature, there is a general gap: the lack of an integrated, adaptable, and operationally useful framework. While there are maturity models and high-level strategies available, organizations still lack one structure that advances strategic aspiration to meet implementation realities. An future-proof framework would need to combine maturity diagnostics with tactical actions, incorporate ethical and feasibility viewpoints, and scale across industries and firm sizes. Only then can GenAI adoption become not just a strategic priority, but an executable advantage.

2.9 Evolution of decision-making frameworks

Over the last several decades, frameworks for decision-making have undergone considerable evolution. Conventional models, including the Rational Decision-Making Model, underscore a systematic and sequential approach wherein decision-makers evaluate various alternatives in pursuit of the best possible solution (Beerbaum, 2023). The bounded rationality theory elucidates the cognitive restrictions faced by human decision-makers, who frequently strive for satisfactory rather than optimal results owing to limitations in information and time.

2.10 AI-integrated decision support systems

Recent advancements have led to AI-integrated decision support systems, such as GenAI Decision Support Systems (GAI-DSS). Chuma et al. (2024) proved that GenAI technologies such as ChatGPT significantly enhance decision-making through processing huge data and offering real-time insights. Studies have shown that AI can enhance organizational efficiency in tasks involving creativity and problem-solving by up to 40%. Even with the capability of these AI systems, however, they tend to lack the infrastructure to offer uniform and effective decision-making outcomes across diverse contexts (Brühl, 2024).

2.11 Structured question-based AI prioritization framework

This paper aims to fill the gap by providing a structured question-based AI framework. It will lead the decision-maker through specific questions to ensure that most aspects of AI are covered and, by doing so, avoid potential mistakes and enhance the quality of the decisions being made. Whereas AI intelligence without structure, think ChatGPT, will have results dependent on the context in which a question is placed, this framework offers a reliable, structured approach. The inquiry-based system supports decision-makers in making full use of the benefits offered by AI while maintaining transparency and systematic rigor in the decision-making process (Candelon, Reeves and Schwarz, 2023).

2.12 Conclusion

The existing literature provides practical recommendations and valuable theoretical insights. However, a critical gap still remains with respect to a structured, question-based system which can guide organization in prioritizing as well as evaluating AI uses cases. As seen in the literature review, most current approaches offer limited support for context-sensitive decision-making, since they e.g. assume high organizational readiness and rely on generic maturity models. The proposed, inquiry-driven, framework in this paper is designed to improve the practical utility, transparency and strategic alignments of GenAI implementations.

3 Methodology

To evaluate GenAI use cases in a way that is consistent, practical, and easy to apply across teams, we have built a structured scoring method. It looks at each use case through three lenses, strategic alignment, technical feasibility, and risk exposure, using a mix of scoring gates, visual tools, and ROI logic. The goal is to move beyond subjective opinions and give teams a repeatable way to assess what is worth investing in, what is scalable, and where the risks really lie. The following sections explain the evaluation model in detail, including how use cases are scored, compared, and prioritized.

3.1 Research approach

This study adopts a design science methodology, focused on creating a practical and adaptable framework, GAIQ (GenAI Qualification), to help organizations evaluate GenAI use cases with clarity and strategic alignment.

Rather than collecting data from individuals, the framework was built through a rigorous review of academic literature, industry reports, and expert consultations. The literature reviews provided insights highlighting the common challenges and considerations in implementing GenAI. Additionally, we engaged with industry experts and practitioners who offered valuable insights based on their real-world experiences. Their feedback helped refine the questions and their relevancy for a Gate, ensuring they addressed practical issues faced by organizations.

This ensured that the model is both theoretically sound and practically applicable in real-world enterprise settings.

3.2 Development of the GAIQ framework

The GAIQ framework was developed to address the growing need for structured evaluation of GenAI initiatives. It is a multi-dimensional, weighted scoring model that integrates business, technical, and ethical parameters.

The framework guides organizations through a gated decision process, helping them assess feasibility, strategic value, and risk before committing resources. By breaking down the evaluation process into these gates and designing targeted questions for each stage, we ensure a comprehensive and structured approach.

This method not only helps in identifying the most promising use cases but also ensures that they are thoroughly vetted before implementation. The theoretical design, grounded in literature and expert insights, provides a robust framework for evaluating GenAI use cases effectively.

GAIQ is designed to be modular and scalable, allowing for customization across industries and organizational maturity levels. It incorporates three core dimensions, PVI (Purpose, Value, and Impact), TFR (Technology Fitment and Resilience), and ERC (Ethics, Risk, and Compliance), which are evaluated across three gates: Scan, Evaluate, and Activate (SEA).

3.3 Definition of evaluation dimensions

The framework evaluates GenAI use cases across three key dimensions:

- PVI: Assesses strategic alignment, measurable outcomes, and operational relevance.
- TFR: Evaluates technical feasibility, integration readiness, and scalability.
- ERC: Reviews ethical implications, data privacy, and regulatory adherence.

Each dimension is weighed differently across the SEA gates to reflect its importance at different stages of evaluation.

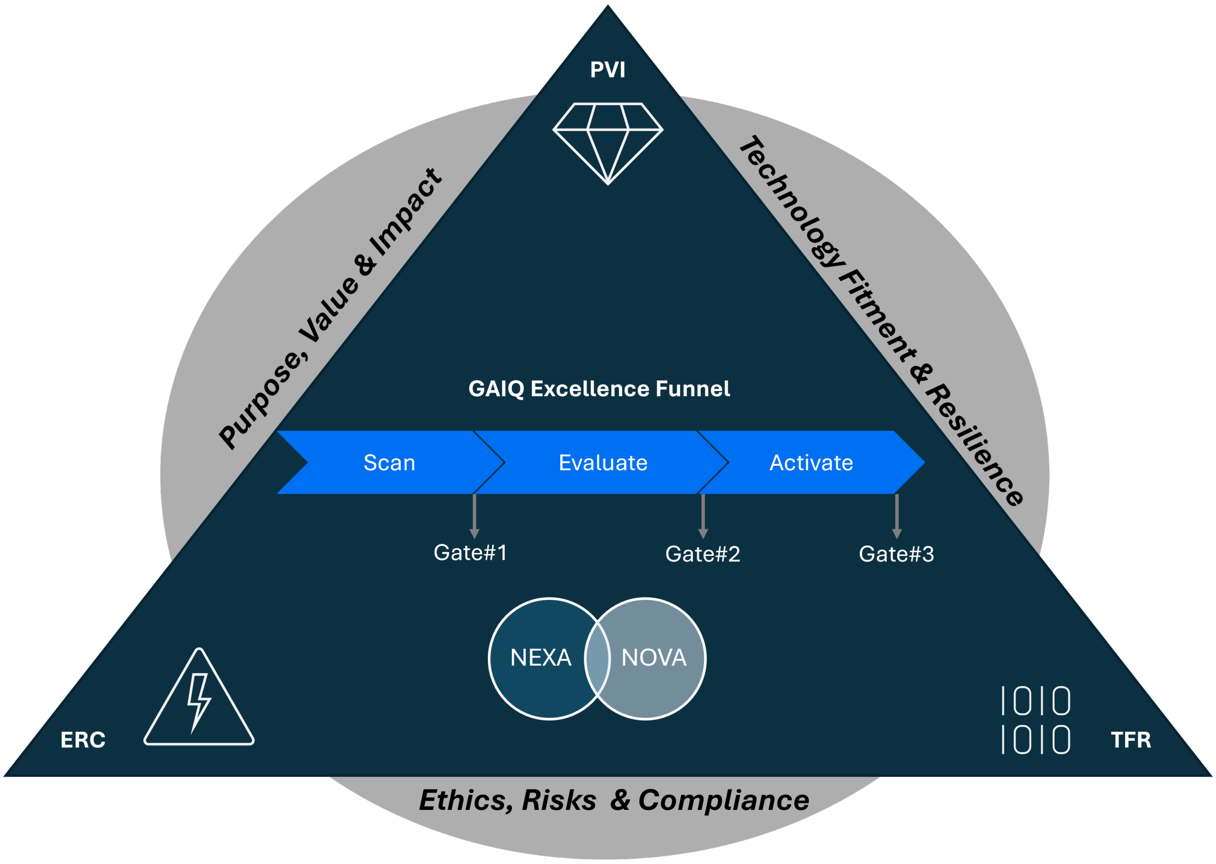

Figure 2. Framework overview

PVI ensures alignment with strategic goals, measurable outcomes, and operational improvements. This involves assessing how the use case supports business objectives, enhances efficiency, and drives innovation. Critical aspects evaluated include:

- Does the AI use case have a clear business purpose?
- Does it align with company priorities?

TFR examines the technical feasibility and resilience of the AI solution. This includes evaluating whether the technology fits within the existing IT infrastructure, its scalability, and its ability to adapt to future changes. It also considers the robustness of the solution in handling various operational scenarios and its resilience against potential disruptions. Key considerations for evaluation include:

- Is the AI technology mature enough?
- Do we have the right data, models, and integrations?

ERC addresses the ethical implications, potential risks, and compliance requirements associated with implementing GenAI. ERC evaluates the ethical implications, potential risks like bias, data sovereignty, algorithmic transparency, and compliance requirements of implementing GenAI. It ensures alignment with legal regulations, safeguards data privacy, and mitigates challenges related to bias, transparency, and accountability. By addressing these factors, ERC supports responsible and trustworthy AI adoption. Key factors evaluated include:

- Does the AI use case meet privacy, security, and legal standards?
- Are biases and governance addressed?

By integrating these three elements, the framework ensures a thorough evaluation, enabling businesses to adopt the most effective and responsible GenAI use cases.

Additionally, the framework also addresses business risk & viability through Next-Gen AI Excellence & Adoption (NEXA) and financial risk through Net Opportunity and Value Assessment (NOVA) adjustment factors to provide a holistic view of feasibility and impact.

3.4 Design of the SEA funnel

The SEA Funnel is a three-stage gated evaluation process. Each gate assesses the use case against a set of questions across three dimensions: PVI, TFR, and ERC. Each use case passes through all Gates sequentially. If it fails one, it does not proceed further.

Scan (Gate 1) focuses on strategic alignment and ethical readiness. In this initial gate, the focus is on identifying potential use cases. The questions here are crucial and comprehensive, aiming to set a high bar for moving forward. This stage emphasizes Purpose, Value, and Impact to Business. The goal is to ensure that only the most promising ideas that align with strategic objectives and ethical standards proceed to the next stage.

Evaluate (Gate 2) assesses technical feasibility and integration complexity. Once the use cases pass through Scan, they move to the Evaluation gate. Here, the questions become more detailed, assessing each use case across all three dimensions but with a balanced distribution. This thorough evaluation helps in understanding the feasibility, potential benefits, and risks associated with each use case.

Activate (Gate 3) prioritizes implementation readiness. The final gate focuses on prioritizing the use cases for implementation. The questions in this stage are designed to ensure that the selected use cases align with the organization's strategic goals and are ready for execution. This stage ensures that only the most impactful and technically feasible projects are chosen. The following figure illustrates the GAIQ Excellence Funnel.

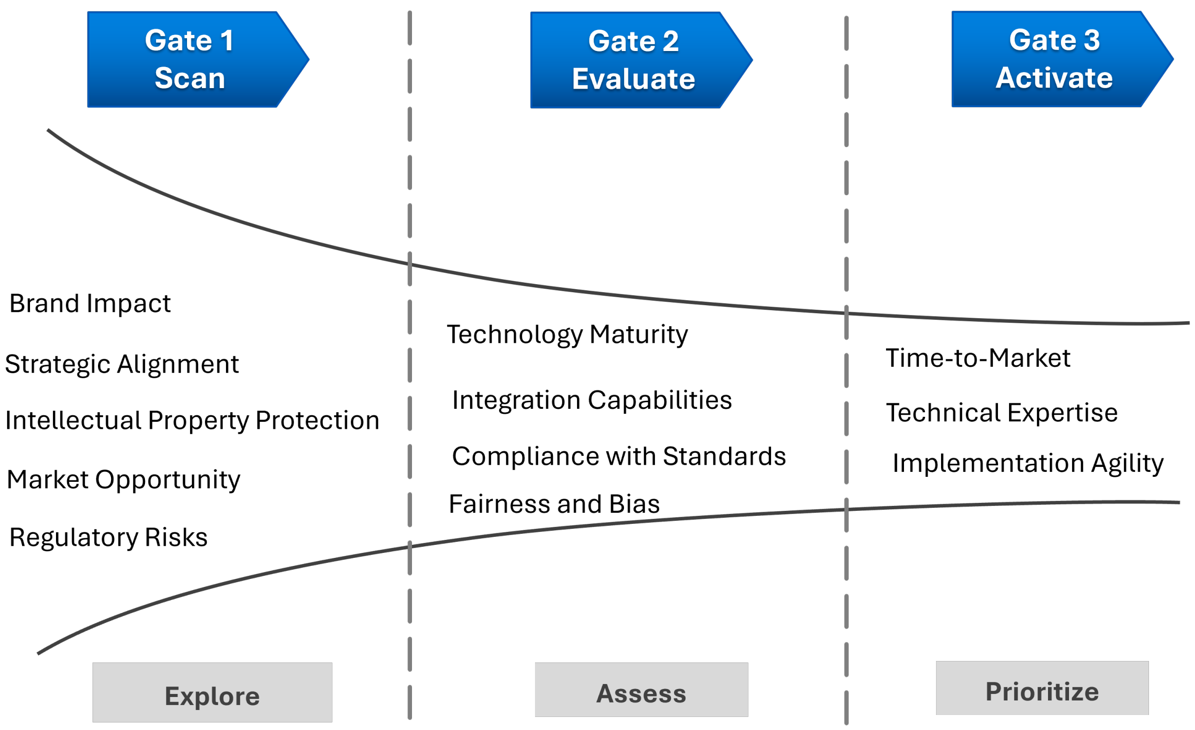

Figure 3.GAIQ excellence funnel

3.5 Scoring and threshold logic

The scoring system uses weighted averages to reflect the relative importance of each dimension. Thresholds are applied at both gate and sub-category levels:

• Gate Threshold: Minimum 75% overall score. If the aggregated score for PVI, TFR, and ERC exceeds this threshold, the use case progresses to the next gate.

• Sub-category Thresholds: Minimum 70% for PVI, TFR, and ERC. For a use case to progress, not only should the aggregated score exceed the gate-level threshold, but each individual category (PVI, TFR, ERC) should also meet the respective threshold.

Use cases must meet both to progress, ensuring balanced evaluation and avoiding advancement of initiatives with critical gaps. These thresholds serve as benchmarks that determine whether a use case is eligible to progress to the next stage. The criteria are designed to maintain balance across key evaluation metrics while identifying areas that may require corrective actions. Below are the two key threshold mechanisms: If a use case fails any threshold, it is flagged for remediation before proceeding. This ensures that no critical dimension is overlooked and promotes continuous improvement. A conditional review mechanism has been added to allow business judgment in borderline cases, enabling revalidation and exception handling.

The model ensures that all important areas (PVI, TFR, and ERC) are thoroughly checked and balanced. This helps avoid pushing forward use cases that might be strong in one area but weak in others, which could cause issues later. It is flexible and adaptable, ensuring focus on areas that need more attention while keeping a thorough review process for each sub-category. With specific thresholds set for both the overall Gate and sub-categories, decisions are made transparently and consistently.

Moreover, the model promotes improvement by highlighting that if a category falls below the threshold, corrective action is needed before moving forward. This ensures continuous improvement. As the number of use cases grows, this model ensures each one is evaluated consistently, with the flexibility to adjust thresholds over time based on past patterns and needs. Having separate gates with their own thresholds makes the decision process clear. At any moment, it is easy to envision why a use case is being allowed or prohibited to progress. This openness encourages accountability and ensures the model is applied uniformly across use cases at all times.

Through the model, a weighted scoring method is employed to determine that the final evaluation accurately depicts the importance of each category. In real scenarios, all the factors do not contribute in equal measures toward the final outcome; certain features are of more importance than others. Assigning weights to various categories allows us to prioritize the most important factors so they contribute more significantly to the final score. This is especially helpful when analyzing systems with more than one criterion that needs to be accounted for, which permits a more detailed analysis than a mere average.

Assigning weights aids in prioritizing categories by importance and relevance. With case, those that have a more direct influence on the outcome are given higher weights so that they can have a stronger influence on the overall score. This approach provides a clearer and more meaningful assessment of performance, one that best shows the relative significance of every factor. While the unweighted mean can be lower in some cases, the weighted score provides a better-balanced evaluation, one that reflects the key factors that lead to success. This approach enhances validity and reliability of our results by anchoring measurement to real-world priorities and ideally capturing relative importance of different criteria.

Briefly, this approach presents a transparent, weighted scoring model with easy decision criteria, easy to use, flexible dimension and gate aggregation and granular insights at question and subcategory levels and a validation process to handle poor scores.. This ensures a fair, robust, and scalable model for a structured evaluation across multiple gates with actionable outcomes for targeted improvement.

3.6 Construction and validation of evaluation questions

The GAIQ evaluation questions were developed through a combination of literature review and expert feedback.

These questions are distributed across the SEA gates and aligned with the three dimensions:

a. Scan: 55% PVI, 40% ERC, 5% TFR. This ensures that the use case is strategically aligned and adheres to necessary privacy and security standards before moving forward.

b. Evaluate: 29% PVI, 42% TFR, 29% ERC. This balanced approach ensures that the use case is not only aligned with strategic goals but also technologically feasible and ethically sound.

c. Activate: 9% PVI, 64% TFR, 27% ERC. This ensures that the use case is ready for deployment and can be executed effectively while maintaining ethical standards and compliance.

Each question is designed to surface critical insights and enable granular scoring. The questions were validated through peer reviews and simulated use cases to ensure relevance and clarity.

3.7 Development evaluation tools

To support structured decision-making, the following tools were developed:

GAIQ – Decision Grid

The decision grid is a visual tool designed to plot and evaluate GenAI use cases based on three critical dimensions: PVI, TFR, and ERC. This matrix is represented as a 9-quadrant map where:

- The X-axis represents the TFR score, ranging from Low to High.
- The Y-axis represents the PVI score, ranging from Low to High.

The size of the bubble represents the ERC score, with larger bubbles indicating higher compliance and lower risk.

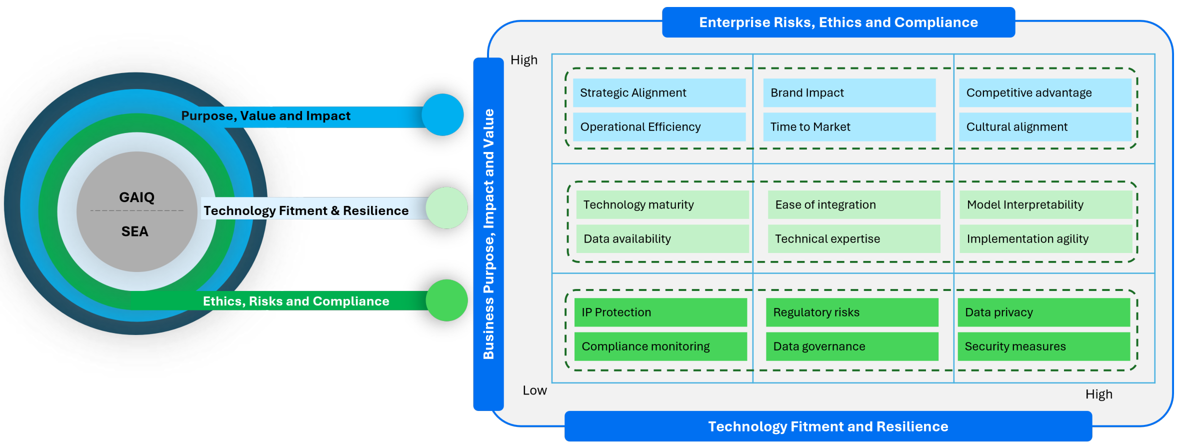

Figure 4. Grid matrix dimensions.

GAIQ Pulse

The GAIQ Pulse is a visual tool designed to plot and evaluate GenAI use cases based on scores of assessments. Leveraging a radar chart format, it plots the composite scores across the three core dimensions. This radar chart provides a clear and insightful representation of each use case's strengths and weaknesses, helping stakeholders make informed decisions based on a thorough understanding of all relevant factors. The illustration below provides an overview of the scoring plotted in the charts.

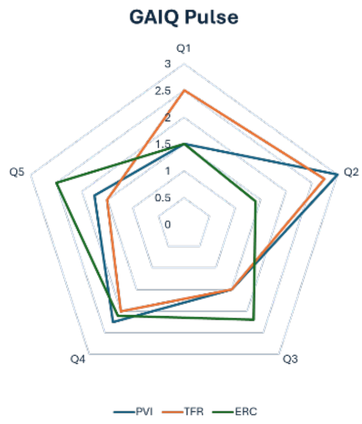

Figure 5. GAIQ pulse

Next Gen AI Excellence and Adoption - NEXA

NEXA is a structured measure that evaluates the viability of AI adoption by balancing AI Enterprise Readiness & Impact Index (AERIX) with ethical, regulatory, and compliance risks i.e. ERC. It ensures that AI projects are not only technologically feasible but also aligned with governance and compliance standards. NEXA also helps businesses understand their strengths and weaknesses, providing a clear picture of where they stand before making big investments in AI. Beyond risk mitigation, NEXA enhances strategic decision-making by helping organizations prioritize AI initiatives based on feasibility and impact. NEXA fosters a balanced approach to AI implementation by aligning technological potential with business priorities, regulatory requirements, and ethical considerations.

AI-driven innovations must not only deliver financial benefits but also comply with evolving regulations and ethical standards. When evaluating a use case, a Risk Penalty Factor (RPF) for ERC risks is considered, which can affect the feasibility of the AI solution. By incorporating these risk factors into AI feasibility assessments, NEXA offers a comprehensive, governance-driven approach to AI adoption, ensuring responsible and sustainable implementation. It acts as a check point, helping decision-makers focus on AI initiatives that they are prepared for, rather than chasing trends.

The AERIX is a composite score that evaluates an organization's readiness for AI adoption. It considers multiple dimensions such as technological infrastructure, data maturity, AI talent availability, and strategic alignment to determine whether an enterprise is ready to implement AI solutions effectively. The AERIX score is calculated by taking the weights assigned to the PVI, TFR, and ERC and multiplying them with the combined PVI, TFR, and ERC scores across all Gates. Essentially, AERIX reflects AI maturity, adoption feasibility, and business impact.

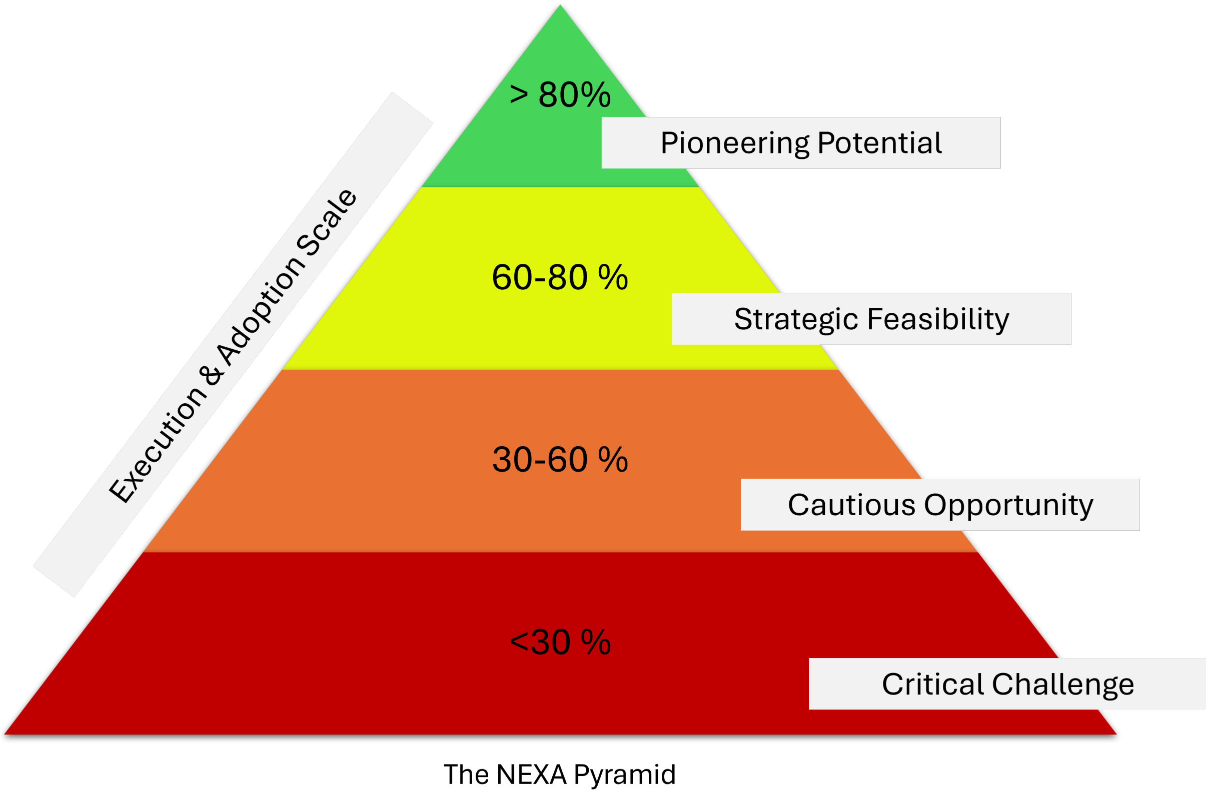

Figure 6. NEXA Pyramid

The table below provides a clear and structured way to understand and make decisions based on the assessment results. It outlines various scenarios and corresponding actions, helping to ensure consistent and informed decision-making. By using this matrix, stakeholders can easily interpret the NEXA results and determine the appropriate next steps.

Tiers	NEXA	Interpretation
Pioneering Potential	>80	AI use case is highly feasible with significant innovation potential, ready for full-scale adoption.
Strategic Feasibility	60-80	Feasible with some risk factors; requires careful planning and risk mitigation before full-scale implementation.
Cautious Opportunity	30-60	Feasible with considerable risks that must be addressed through detailed mitigation strategies.
Critical Challenge	<30	High-risk, low feasibility use case. Requires major improvements or complete re-evaluation before proceeding

Table 2. Interpretation and decision matrix.

Net Opportunity and Value Assessment - NOVA

The NEXA score is derived from AERIX and adjusted by a Risk Penalty Factor. However, it does not incorporate the financial viability of the AI use case nor the risks that are unique for the industry.

NOVA is designed to quantify and manage financial risks associated with AI investments. At its core, NOVA integrates AI Investment Risk-Adjustment (AIRAx), a methodology that evaluates both upfront capital exposure and long-term financial viability of AI projects. Traditional ROI models often overlook hidden cost uncertainties, regulatory liabilities, and scalability constraints. NOVA bridges this gap by incorporating AIRAx as a weighted factor in decision-making. Traditional ROI would have been much higher, but after adjusting for AI-specific risks, the realistic financial outcome is far lower. For example, AI in manufacturing has high operational risks, so true TCO must factor in downtime, integration costs, and scalability risks.

Additionally, Industry Risk Multiplier (IRM) is also considered to adjust the risk based on varying levels of AI adoption maturity, regulatory complexity across different industries. Industries with high compliance risk (e.g., healthcare, finance) will have an IRM adjustment bias toward regulatory risks.

While IRM accounts for industry-specific risks, AIRAx adjusts for financial volatility, economic conditions, and AI adoption uncertainty at an organizational level. This methodology helps in preventing underestimation of AI implementation risks in financial models and ensuring projects are technical feasible, ethically sound and financially viable. The following illustration summarizes the NOVA framework and its components.

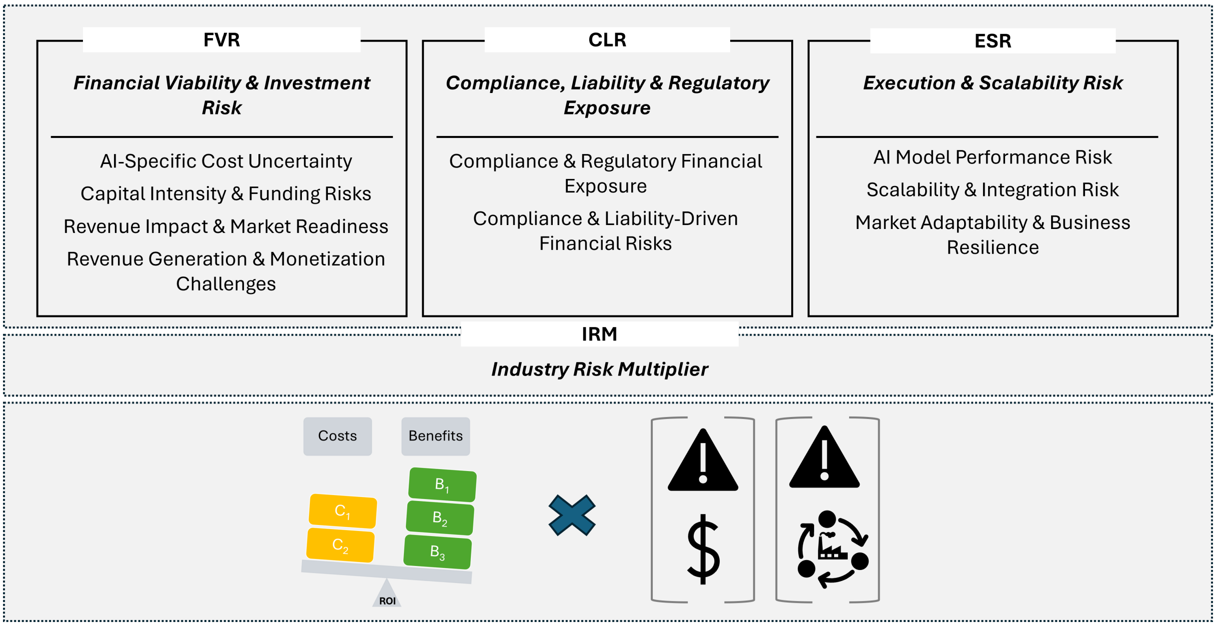

Figure 7. NOVA framework

There are three broad dimensions in arriving at the AIRAx factor:

- FVR (Financial Viability & Investment Risk): This risk involves the potential for AI investments to be financially unsound, leading to wasted resources and poor returns on investment.
- CLR (Compliance, Liability & Regulatory Exposure): This risk pertains to the possibility of AI adoption resulting in legal issues or financial liabilities due to non-compliance with regulations or unforeseen legal challenges.
- ESR (Execution & Scalability Risk): This risk concerns the stability, scalability, and adaptability of AI solutions, which may fail to perform as expected under different conditions or when scaled up, leading to operational disruptions and inefficiencies.

Sub-Parameter	Criteria	Score Range	Description
FVR	ROI potential, cost efficiency, funding risk	0–3	Higher score = higher financial risk or poor investment viability
CLR	Legal exposure, data privacy, regulatory complexity	0–3	Higher score = greater risk of non-compliance or legal liability
ESR	Technical stability, scalability, adaptability	0–3	Higher score = higher risk of operational failure or poor scalability

Table 3. Dimensions in arriving at the AIRAx-factor

IRM enables Industry-specific adjustments improve financial decision-making accuracy. Different industries face varying levels of regulatory scrutiny, market volatility, and technological adoption barriers, which directly influence AI success rates. IRM quantifies these external uncertainties, acting as a scaling factor that moderates the projected ROI based on industry-wide trends and challenges. IRM adjusts these AI investments primarily based on 3 factors:

- Regulatory Compliance (GDPR, HIPAA, AI Act, Financial compliance etc.)
- Liability risks (AI Bias, Ethical issues, litigation risks)
- Execution & Resilience (Security vulnerabilities and other enterprise risks)

3.8 Validation through simulated use cases

To ensure the GAIQ framework was not just theoretically sound but practically applicable, it was validated using a series of simulated decision scenarios. These simulations were designed to reflect real-world challenges across diverse domains such as customer support, marketing, and code generation, all areas where GenAI is increasingly being adopted. The use cases were evaluated without real enterprise data and analyzed from the perspective of how well they can hold to realistic, business relevant situations.

Each of these use cases were passed through the SEA Gates to assess their alignment with business goals, technical feasibility and scalability and review for bias, privacy and regulatory compliances.

In addition to the SEA gates, the GAIQ pulse was also reviewed to understand the scores across the three core dimensions and interpretation of the results. Following this, the NEXA matrix was generated to identify the potential of the use case and the recommendation Tier. The NOVA was determined by analyzing the FVR, CLR and ESR for the use case to determine the AIRAx score. Lastly, the Adjusted ROI was calculated based on the NOVA and IRM score.

4 Results

To ensure consistent and credible evaluation of GenAI business use cases, we followed a structured flow that begins with strategic alignment and ends with financial viability. This approach helps stakeholders make informed decisions by balancing innovative potential with implementation readiness and financial risk. The following section summarizes the three use cases considered for evaluating the model and presents the results from the Funnel to Risk-adjusted ROI.

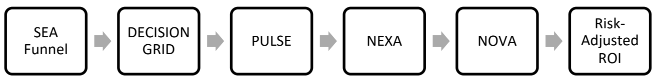

Figure 8. Results from funnel to risk-adjusted ROI.

4.1 Description of simulated use cases

The GAIQ framework was applied to three hypothetical use cases:

- Use case 1 - Agentic Customer Support: An AI agent designed to automate and enhance customer service.
- Use case 2 - AI-driven Market Analysis Tool: A solution that provides real-time market insights for strategic planning.
- Use case 3 - Personalized Marketing Generator: A GenAI engine that creates tailored marketing campaigns based on customer data.

Gate Evaluation Outcomes for the three different use cases are illustrated in tables 4-6.

Gate	PVI	TFR	ERC	Overall	Comments
Gate 1	78	80	77	78	All scores above threshold; strong business alignment and technical feasibility.
Gate 2	76	78	72	75	ERC below threshold due to legal complexity and compliance risks.
Gate 3	74	76	70	73	ERC still below 70%; requires validation of regulatory safeguards.
Final	**76**	**78**	**73**	**76**	Final score below 75%; proceed with risk mitigation plan.

Table 4. Use case 1: Agentic customer support

Gate	PVI	TFR	ERC	Overall	Comments
Gate 1	85	88	40	71	ERC significantly below threshold due to ethical concerns and brand safety.
Gate 2	65	82	45	64	PVI and ERC below threshold; reassess strategic value and compliance.
Gate 3	68	80	50	66	ERC remains low; requires governance and ethical content controls.
Final	**73**	**83**	**45**	**67**	Final score below 75%; use case not ready for scaling.

Table 5. Use case 2: AI-driven market analysis

Gate	PVI	TFR	ERC	Overall	Comments
Gate 1	90	70	50	70	TFR and ERC are significantly below threshold due to integration challenges and ethical concerns around personalization and data privacy.
Gate 2	78	65	60	68	TFR and ERC remain below threshold; strategic value is strong but compliance and technical readiness require further validation.
Gate 3	80	70	65	72	ERC still below threshold; governance and ethical safeguards must be addressed before scaling.
Final	**83**	**68**	**58**	**70**	Final score below 75%; use case not ready for deployment without additional compliance and oversight measures.

Table 6. Use case 3: Personalized Marketing Content

To ensure that promising use cases are not prematurely disqualified, the following approach is recommended:

a) assess whether the overall score shortfall is due to business feasibility constraints (PVI) or a temporary misalignment in weighting factors.
b) revalidating the business impact with domain experts and business sponsors and decision-makers can apply business judgment instead of a strict numeric cutoff.
c) identify opportunities for re-assessment where targeted improvements are made, and business case is re-submitted.

d) conditional reviews and sign-off during the next Gate or if decision is at the end of the cycle, the use case should not proceed without significant compliance and agreement on formal metrics to evaluate success and risks.

4.2 Pulse radar chart results

The illustrations below provide an overview of the scoring plotted in the charts as an example to demonstrate the details available for deep dive and to take necessary course correction while the use case is under review.

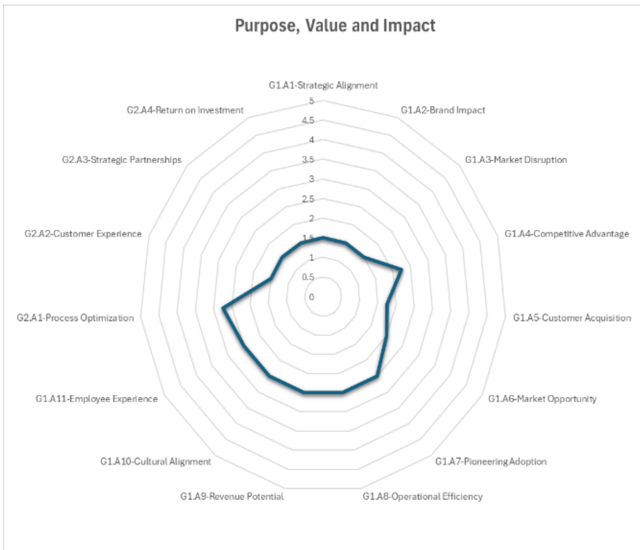

Figure 9. PVI pulse.

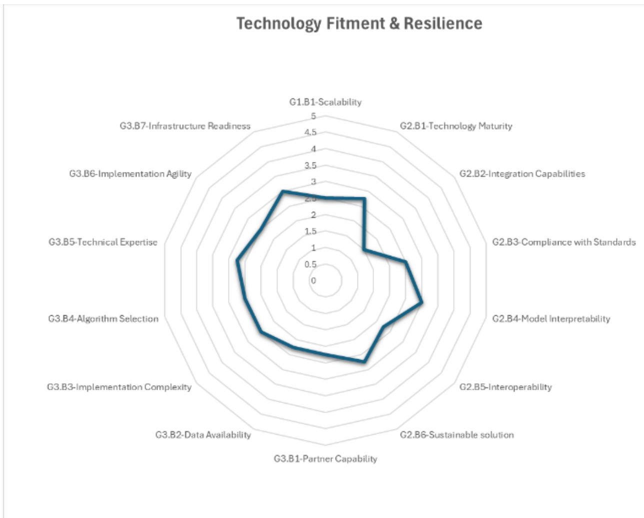

Figure 10. TFR pulse.

Figure 11. ERC pulse.

By analysing the charts, stakeholders can quickly identify areas where a use case excels or needs improvement. They can leverage the numerical scores to highlight areas that need to be addressed before implementation, such as improving technical integration or ensuring better compliance with regulations. Accurate evaluation is critical to keep scoring realistic and unbiased.

The GenAI domain is fast evolving, expanding the opportunities for applications in the industry. A continuous review and updates to the evaluation criteria and scoring based on the new information and feedback ensures that the framework remains relevant and accurate. This structured approach helps in maximising the benefits of AI technologies while mitigating potential risks and challenges.

4.3 Decision grid analysis Implications

The chart below summarizes the results of the three sample scenarios explained above. The decision provides a comprehensive way to visualize and prioritize GenAI use cases. By plotting the PVI, TFR, and ERC scores, businesses can easily identify which use cases offer the most value, are technically feasible, and comply with ethical and regulatory standards. This structured approach ensures that the most impactful and responsible AI projects are selected for implementation.

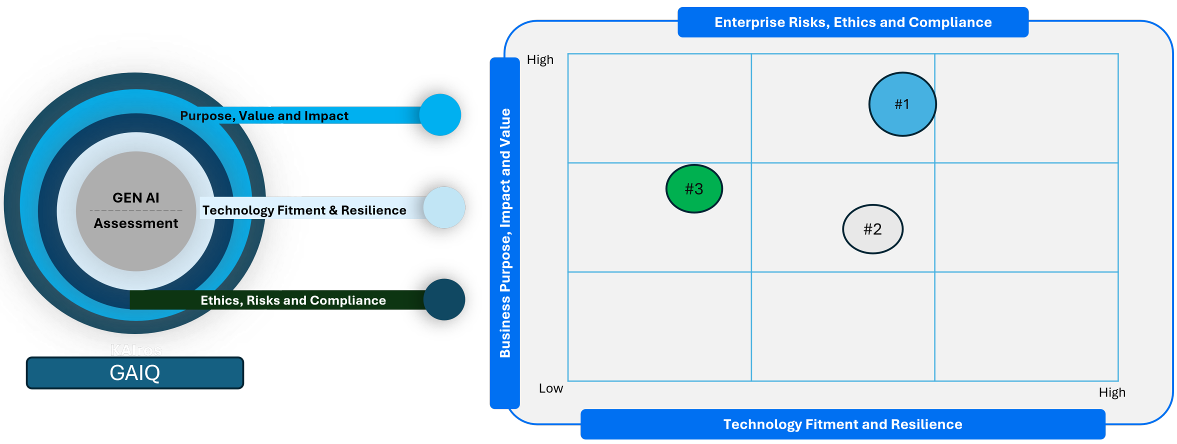

Figure 12. Example of grid matrix.

4.4 NEXA and NOVA scoring outcomes

The following tables illustrates the NEXA calculation for the three use cases, including recommendations. It is assumed that the weights for PVI, TFR and ERC are (0.4, 0.3, 0.3).

Step 1 - Calculate Aggregate Implementation Readiness Index (AERIX)

$$AERIX = (0.4 \times PVI) + (0.3 \times TFR) + (0.3 \times ERC)$$

Step 2 - Compute Risk Penalty Factor (RPF)

$$RPF = 1 + \left(\frac{100 - ERC}{100}\right)$$

Step 3 - Compute NEXA

$$NEXA = \frac{AERIX}{RPF}$$

Calculation	Use Case 1	Use Case 2	Use Case 3
AERIX	76	68	71
RPF	1.27	1.55	1.42
NEXA	60	44	50

Table 7. NEXA calculations.

Use Case	NEXA	Action / Recommendation
Use case 1: Agentic customer support	**60**	Strategic Feasibility
Use case 2: AI-driven market analysis	**44**	Cautious Opportunity
Use case 3: Personalized Marketing Content	**50**	Cautious Opportunity

Table 8. Recommendations as per the NEXA pyramid.

Calculating the NOVA
The formula for calculation is as follows:

$$\text{Adjusted ROI} = \left(\frac{\text{Projected ROI}}{\text{AIRAx} \times \text{IRM}}\right) \times 100$$

Where AIRAx is calculated as

$$AIRAx = (w_1 \times FVR + w_2 \times CLR + w_3 \times ESR) \div (w_1 + w_2 + w_3)$$

Where w_1, w_2, w_3 = weights assigned to each risk dimension based on business impact

The simulation example is presented for a Consumer goods industry, applicable IRM has been considered for the calculations. A financial summary of the use cases is presented in below table.

Calculation	Use Case 1	Use Case 2	Use Case 3
1. CAPEX	$ 1.50 Mn	$ 1.25 Mn	$ 2.5 Mn
2. OPEX / year	$ 500 K	$ 400 K	$ 600 K
3. OPEX / 3 year	$ 1.5 Mn	$ 1.2 Mn	$ 1.8 Mn
4. Projected Return (Top Line & Bottom-line impact) – 3 Year total	$ 6 Mn	$ 4 Mn	$ 10 Mn
5. AIRAx	(0.5*2+0.4*2+0.3*1.5) / (0.5+0.4+03) = 1.875	(0.5*2.5+0.4*1.75+0.3*2) / (0.5+0.4+03) = 2.125	(0.5*2+0.4*2+0.3*2) / (0.5+0.4+03) = 2
6. IRM	1.3	1.3	1.3
7. ROI	**200%** [6-(1.5+1.5)] /1.5	**125%** [4 – (1.25+1.2)] /1.25	**317%** [10-(2.5+1.8)] /2.5
8. Adjusted ROI	**107%** 200% /(1.875*1.3)	**59%** 125% / (2.125*1.3)	**132%** 317% / (2*1.3)

Table 9. A financial summary of the use cases. All values in Mn USD.

In summary the NOVA enables a CXO to arrive at a credible risk-adjusted ROI making investment decisions financially sound and industry-aware.

4.5 Key evaluation patterns and trade-offs

The evaluation of use cases across different domains revealed several patterns and insights:

- High PVI Scores: Use cases with high PVI scores were often those that directly aligned with strategic business objectives and had clear, measurable outcomes.
- Technical Feasibility: Use cases with high TFR scores were those that could be easily integrated with existing systems and had mature AI technologies.
- Ethical Considerations: Use cases with high ERC scores were those that adhered to regulatory standards and addressed ethical concerns such as data privacy and bias.

Anomalies and Implications
While most use cases followed the expected patterns, some anomalies were observed:

- High PVI, Low TFR: Some use cases had high strategic value but faced technical challenges, indicating the need for further investment in technology and infrastructure.
- High TFR, Low ERC: Use cases that were technically feasible but had ethical or regulatory concerns highlighted the importance of addressing these issues early in the evaluation process.

Framework Utility

- Granular Scoring: The SEA model's gate-wise weighting exposed early-stage weaknesses Objective Thresholds: The 75% gate threshold ensured only well-rounded use cases progressed.
- Radar Chart Logic: Enabled visual identification of imbalances (e.g., strong TFR but weak ERC).

The results demonstrated the framework's ability to:

- Identify weak areas (e.g., low ERC in marketing AI).
- Prevent premature advancement of high-risk use cases.
- Provide actionable insights for improvement and re-evaluation.

4.6 Comparison with traditional evaluation methods

While generic AI maturity models and strategic frameworks from Gartner, McKinsey, and PwC offer valuable top-down guidance for enterprise-wide AI adoption, they often lack a repeatable, bottom-up methodology for evaluating individual AI use cases. These models typically focus on organizational readiness, transformation roadmaps, or value loops, but do not provide the granular, scenario-specific decision support needed for operational prioritization.

The GAIQ framework addresses this critical gap by offering:

- A structured, gate-based evaluation model (Scan–Evaluate–Activate) that can be applied consistently across diverse AI initiatives.
- Quantitative scoring across 46 sub-dimensions spanning business value (PVI), technical feasibility (TFR), and ethical compliance (ERC).
- Visual and decision-support tools such as radar charts, decision grids, and readiness indices (NEXA, NOVA) to guide prioritization.
- A flexible, modular design that supports domain-specific calibration and stakeholder alignment.

This makes GAIQ uniquely suited for organizations seeking to scale GenAI responsibly, with precision and strategic alignment at the use case level. This addresses the key gaps typically in existing models like:

- Granularity: Lacking detailed, scenario-specific scoring
- Repeatability: Subjective assessments vary across reviewers
- Operationalization: Limited support for visual tools or integration into agile workflows

Its structured approach simplifies decision-making, provides deeper insights, and ensures that organizations focus on high-impact, strategically aligned, and ethically sound AI initiatives.

5 Discussion

This discussion interprets GAIQ's practical implications, adoption hurdles, limitations, and next research steps.

5.1 Implications

The GAIQ framework is proving to be a game-changer for businesses trying to make sense of GenAI opportunities. It is like having a GPS for navigating AI decisions, helping companies zero in on the projects that will really move the needle. For highly regulated industries like banking and healthcare, the ERC component is particularly valuable, keeping companies on the right side of new regulations like the EU AI Act. Meanwhile, retail and manufacturing companies are finding the PVI and TFR aspects especially useful for spotting opportunities that can drive real growth and streamline operations.

5.2 Limitations

While the GAIQ framework offers a robust approach to evaluating GenAI use cases, it has certain limitations. The framework has been validated through simulated scenarios and peer feedback, but it has not yet been tested at scale across multiple organizations. The scoring system is solid, but it probably needs fine-tuning to better reflect different industry realities. Moving forward, we need to put the framework through its paces across various industries and refine how we measure success to make it even more accurate and useful.

5.3 Stakeholder engagement and governance

GAIQ is best operationalized when decision rights are explicit (see Table 10). Delegate ownership by dimension, Business → PVI, Technology → TFR, Risk/Legal → ERC, and move each use case through three gates: Scan, Evaluate, and Activate.

In Scan, initial fit and risk are approved by the business sponsor or domain lead, with co-signed by compliance and input from enterprise/cloud architects; critical materials include the problem statement, value hypothesis, data provenance, and pre-assessment. In Evaluate, data science/engineering lead on TFR, co-signed by product and finance for PVI and compliance/security for ERC. Materials include the integration checklist, model card, privacy/security reviews, and NOVA sheet. In Activate, engineering/MLOps lead deployment readiness with co-sign-off from the business (PVI KPIs/benefits) and risk/legal/security (ERC controls). Materials include a deployment plan, monitoring KPIs, model-risk controls, and a rollback plan. This governance reduces reviewer variability, ties controls to risk, and improves auditability.

Category	Role	PVI	TFR	ERC
Business	Business Sponsor	●	◑	◑
	Domain Experts	●	◑	○
	Financial Analysts	◑	●	○
	Product Managers	●	◑	◑
	Innovation Champions	●	◑	◑
	Change Specialists	◑	●	◑
Technology	Data Scientists	◑	●	◑
	AI/ML Engineers	◑	●	◑
	Enterprise Architects	◑	●	◑
	Cloud Architects	◑	●	◑
	Technology Advisors	◑	◑	◑
	Data Engineers	◑	●	◑
Compliance	Compliance Officers	◑	◑	●
	Security Specialists	◑	●	●
	Risk Analysts	●	◑	●

Table 10. Stakeholder matrix

6 Conclusion

GAIQ delivers a repeatable, gate-based method to select high-value, low-risk GenAI use cases. Calibrated PVI/TFR/ERC scoring, complemented by NEXA and NOVA, reduces subjectivity and captures ERC/TFR gaps early. It also enables CFO-grade investment decisions.

Deployed in three virtual business case environments, GAIQ identified initial ERC/TFR risk areas, prevented premature advancement of high-risk projects, and enabled transparent portfolio sequencing. Numerical scoring (46 sub-dimensions) and graphics supported targeted remediation and stakeholder alignment.

This paper has two boundary conditions: experimentations with used simulations against expert judgment rather than multi-firm field data, and weights/thresholds were pragmatically adjusted and require sector-specific refinement. GAIQ can be used for selection and ordering; post-deployment surveillance and model-risk controls should supplement it in practice.

Future research needs to (i) field studies comparing GAIQ with unstructured approaches in terms of time-to-value, risk incidents, and ROI; (ii) reliability test (multi-rater agreement, sensitivity of NEXA/NOVA to ERC attenuation); (iii) prototype first-pass scoring LLM-assisted; and (iv) extend the rubric with sustainability and environmental impact measurement.

References

Amazon Web Services (2025) 'What is generative AI?'. Available at: https://aws.amazon.com/what-is/generative-ai/ (Accessed: 5 June 2025).

Baxter, K. and Schlesinger, Y. (2023) 'Managing the risks of generative AI', *Harvard Business Review*, 6 June. Available at: https://hbr.org/2023/06/managing-the-risks-of-generative-ai (Accessed: 6 September 2025).

Beerbaum, D. (2023) 'Generative artificial intelligence (GAI) with ChatGPT for accounting - a business case', *SSRN Electronic Journal*, 11 March. Available at: https://ssrn.com/abstract=4385651 (Accessed: 6 September 2025).

Brühl, V. (2024) 'Generative artificial intelligence - foundations, use cases and economic potential', *Intereconomics*, 59(1), pp. 5–9. https://doi.org/10.2478/ie-2024-0003.

Candelon, F., Reeves, M. and Schwarz, J. (2023) *Generative AI decision-making frameworks*. Boston, MA: Boston Consulting Group. Available at: https://www.bcg.com/publications/generative-ai-frameworks (Accessed: 6 September 2025).

Challapally, A., Pease, C., Raskar, R. and Chari, P. (2025) *The GenAI divide: State of AI in business 2025*. Cambridge, MA: Project NANDA. Available at: https://www.artificialintelligence-news.com/wp-content/uploads/2025/08/ai_report_2025.pdf (Accessed: 13 September 2025).

Chuma, E.L., Alves, A.M. and de Oliveira, G.G. (2024) 'Evolution of generative AI for business decision-making: A case of ChatGPT', *Management Science and Business Decisions*, 4(1), pp. 5–14. https://doi.org/10.52812/msbd.87.

Deloitte (n.d.) 'Deloitte's Trustworthy AI™ framework and the White House Blueprint for an AI Bill of Rights'. Available at: https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/articles/ai-bill-of-rights.html (Accessed: 6 September 2025).

Deloitte AI Institute (2024) *The state of generative AI in the enterprise 2024*. New York: Deloitte Insights. Available at: https://www2.deloitte.com/us/en/pages/consulting/articles/state-of-generative-ai-in-enterprise.html (Accessed: 6 September 2025).

Ebers, M. (2024) 'Truly risk-based regulation of artificial intelligence: How to implement the EU's AI Act', European Journal of Risk Regulation, pp. 1–20. Available at: https://doi.org/10.1017/err.2024.78.

Ernst & Young (2024a) 'Artificial-intelligence investments set to remain strong in 2025, but senior leaders recognize emerging risks', *EY Newsroom*, December. Available at: https://www.ey.com/en_us/newsroom/2024/12/ey-research-artificial-intelligence-investments-set-to-remain-strong-in-2025-but-senior-leaders-recognize-emerging-risks(Accessed: 6 September 2025).

Ernst & Young (2024b) *EY AI Pulse Survey 2024*. London/Global: Ernst & Young Global. Available at: https://www.ey.com/en_us/insights/emerging-technologies/pulse-ai-survey (Accessed: 6 September 2025).

Gartner (2023) *AI maturity model: A roadmap for AI adoption*. Stamford, CT: Gartner. Available at: https://www.gartner.com/en/chief-information-officer/research/ai-maturity-model-toolkit (Accessed: 6 September 2025).

Greenstein, D., Light, A. and Likens, S. (2024) *PwC's GenAI value-realization flywheel*. London: PwC. Available at: https://www.pwc.com/gx/en/issues/technology/path-to-generative-ai-value.html (Accessed: 6 September 2025).

Gupta, V. (2024) ‘An empirical evaluation of a generative artificial intelligence technology adoption model from entrepreneurs’ perspectives’, *Systems*, 12(3), 103. https://doi.org/10.3390/systems12030103.

IBM (2024) ‘What is generative AI?’, *IBM Think*, 22 March. Available at: https://www.ibm.com/think/topics/generative-ai (Accessed: 6 September 2025).

Kalusivalingam, A.K., Sharma, A., Patel, N. and Singh, V. (2022) ‘Leveraging generative adversarial networks and reinforcement learning for business model innovation: A hybrid approach to AI-driven strategic transformation’, *International Journal of AI and ML*, 3(9). Available at: https://www.cognitivecomputingjournal.com/index.php/IJAIML-V1/article/view/66 (Accessed: 6 September 2025).

Lenovo (2025) ‘Lenovo global study reveals ROI remains greatest AI adoption barrier, despite three-fold spend increase’, *Press release*, 5 February. New York: Lenovo. Available at: https://news.lenovo.com/pressroom/press-releases/cio-playbook-global-study-roi-ai-adoption/ (Accessed: 6 September 2025).

Lindéus, G. and Shetty, S. (2024) “Global AI Readiness Clusters: Patterns and Disparities in Governance and Technological Infrastructure”, *Global journal of Business and Integral Security*, 7(3). Available at: https://www.gbis.ch/index.php/gbis/article/view/600

Master, D., Chirputkar, N. and Ashok, R. (2024) ‘Generative AI in business: A practical guide’, *AI Practice Review*, 8(1), pp. 22–34.

McKinsey & Company (2023) *The state of AI in 2023: Generative AI’s breakout year*. New York: McKinsey & Company. Available at: https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2023-generative-ais-breakout-year (Accessed: 6 September 2025).

McKinsey Global Institute (2023) *The economic potential of generative AI: The next productivity frontier*. New York: McKinsey & Company. Available at: https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier (Accessed: 6 September 2025).

Nguyen, K.V. (2025) ‘The use of generative AI tools in higher education: Ethical and pedagogical principles’, *Journal of Academic Ethics*. Available at: https://link.springer.com/article/10.1007/s10805-025-09607-1 (Accessed: 6 September 2025).

Patil, A., Rane, S. and Rane, S. (2024a) *Maximizing ROI through generative AI applications*. Pune: AI Strategy Press.

Patil, A., Rane, S. and Rane, S. (2024b) *AI impact prioritization for business growth*. Mumbai: TechEdge Publishing.

Patil, A., Rane, S. and Rane, S. (2024c) *Cybersecurity imperatives in generative AI*. Delhi: AI Ethics Institute.

Rajaram, K. and Tinguely, P.N. (2024) ‘Generative artificial intelligence in small and medium enterprises: Navigating its promises and challenges’, *Business Horizons*, 67(5), pp. 629–648. https://doi.org/10.1016/j.bushor.2024.05.008.

Ransbotham, S., Kiron, D., LaFountain, B., Kudaravalli, S., Candelon, F., LaFountain, M. and Khodabandeh, S. (2019) *Winning with AI*. Cambridge, MA: MIT Sloan Management Review and Boston Consulting Group. Available at: https://sloanreview.mit.edu/projects/winning-with-ai/ (Accessed: 6 September 2025).

Salesforce (2023) ‘Top generative AI statistics for 2023’, *Salesforce News & Insights*, 1 September. Available at: https://www.salesforce.com/au/news/stories/generative-ai-statistics/ (Accessed: 6 September 2025).

Samala, A.D., et al. (2024) ‘Unveiling the landscape of generative artificial intelligence in education’, *Education and Information Technologies*. Available at: https://link.springer.com/article/10.1007/s10639-024-12936-0 (Accessed: 6 September 2025).

Sterne, J. (2023) ‘Measuring the business value of generative AI’, *Journal of AI, Robotics & Workplace Automation*, 3(1), pp. 28–36. https://doi.org/10.69554/AUIJ4734.

van der Veen, B. (2024) *Bridging the generative leap: Exploring factors determining generative AI-readiness in organizations*. Master’s thesis. Eindhoven: Eindhoven University of Technology. Available at: https://research.tue.nl/nl/studentTheses/bridging-the-generative-leap (Accessed: 6 September 2025).

Wala, T. and Wooten, E. (2024) ‘Implementing generative AI at large enterprises: The technology is not where big companies struggle’, *Journal of Business and Artificial Intelligence*, 1(1), 16 May. Available at: https://jbai.ai/index.php/jbai/article/view/24 (Accessed: 6 September 2025).

Zao-Sanders, M. (2025) ‘How people are really using Gen AI in 2025’, *Harvard Business Review*, 9 April. Available at: https://hbr.org/2025/04/how-people-are-really-using-gen-ai-in-2025 (Accessed: 6 September 2025).