Everest Group®

November 2024

Assuring the Rewards of Generative AI

Navigating the Future Responsibly

Contents

Executive summary

Enterprises are adopting generative AI use cases across functions, anticipating multiple rewards. These rewards range from efficiency-centric outcomes, such as streamlined operations and enhanced productivity, to growth-centric rewards, such as increased profits and new revenue streams. However, enterprises are exposed to various risks related to data security, compliance, decision biases, and output accuracy.

Our recent survey with global IT and business leaders indicates that risk remains a key consideration when implementing generative AI use cases. Some of our key research findings are:

- Enterprises are currently prioritizing use cases with minimal to limited risk for generative AI, resulting in sparse rewards for them
- Around 85% of enterprises believe that they are aware of generative AI risks but are unprepared to handle them to achieve higher rewards

Enterprises are investing in building robust risk management strategies to maximize current and future rewards from generative AI use cases. However, as per our survey, only two in five enterprises have matured in their generative AI risk management strategies. These enterprises are realizing 25% higher rewards in the form of improved customer experience, productivity, and speed to market.

As we examined these mature enterprises' characteristics, we found:

- These enterprises strongly believe that the quality function will play a significant role in assuring generative AI's rewards for business- and IT-specific use cases
- They have identified the critical role of the quality function in assuring every layer of generative AI applications and the at every stage of the adoption cycle

This report examines potential risks of generative AI adoption and presents a roadmap for managing risks and sustainability scaling adoption across the enterprise. The report also presents a recommended approach for the quality function to assuring the expected rewards from generative AI investments.

Our research uncovered the following key recommendations for enterprises to maximize generative AI rewards:

- **Establish generative AI risk management as a core enterprise agenda:** Enterprises must be aware of the risks they are exposed to in their generative AI adoption journeys. They need to consistently invest in enterprise-wide risk management initiatives across people, processes, and technology levers.
- **Empower the quality function to maximize rewards from generative AI investments:** The quality function is well positioned to advocate for risk management initiatives with the generative AI project execution teams. Training the quality function to take on the evolved responsibilities of assuring variable outcomes of generative AI applications and systems can enable it to become the custodian for managing risks and assuring rewards.
- **Redesign quality metrics to assure enterprise rewards from generative AI:** Enterprises must redefine performance metrics to accommodate the dynamic nature of generative AI ecosystems. Enterprises should establish robust governance to continuously track the metrics essential to maintaining the quality and robustness of generative AI systems. This will enable enterprises to ensure their applications remain resilient, enabling responsible AI adoption and maximizing rewards.

Embracing generative AI's evolution in business

> **Everest Group take**
>
> Enterprises are adopting generative AI use cases to drive efficiency-focused outcomes. However, they recognize that this adoption comes with risks, including data security, bias, accuracy, and fairness of results. As a result, several use cases that enterprises are implementing have low to moderate risk exposure, and commensurate expected rewards. Implementing a well-defined risk management strategy is essential for enterprises to successfully scale these initiatives and unlock generative AI's potential.

Over the past year, enterprises have demonstrated significant enthusiasm for adopting generative AI use cases across processes and functions. Many enterprises are now moving beyond the exploratory or Proof of Concept (PoC) phase and scaling their generative AI initiatives. Additionally, companies allocate dedicated budgets in their IT expenditures for generative AI projects.

1 of 3 enterprises allocate over 5% of their IT budgets for generative AI initiatives.

As enterprises increase their generative AI adoption, their objectives evolve beyond operational efficiency and productivity. In the short term, firms aim to improve processes and reduce time to market. However, as investments scale and mature, they target higher-value returns, such as enhancing decision-making capabilities and fostering business innovation. This shift highlights a deeper understanding of generative AI's potential to drive long-term growth and transform enterprise functions beyond efficiency.

Exhibit 1 lists enterprise expectations from generative AI use cases.

Exhibit 1: Enterprise expectations from the generative AI use cases

Source: Everest Group (2024)

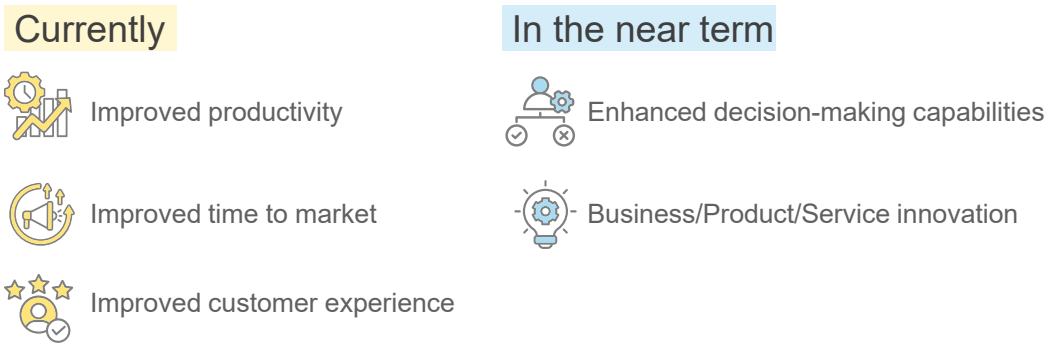

87% of enterprises have budgets for generative AI initiatives allocated centrally as part of the overall planning process.

With these evolving ambitions, the risk landscape also changes. Initially, enterprises prioritized managing data security risks and biases, which are most pressing during the early adoption stages. However, as generative AI initiatives scale and become more integrated into core operations, new risks related to model outcomes, such as accuracy, fairness, and ethical considerations, come to the forefront.

Exhibit 2 shows various risks enterprises encounter while executing generative AI use cases.

Exhibit 2: Risks enterprises encounter while executing generative AI use cases

Source: Everest Group (2024)

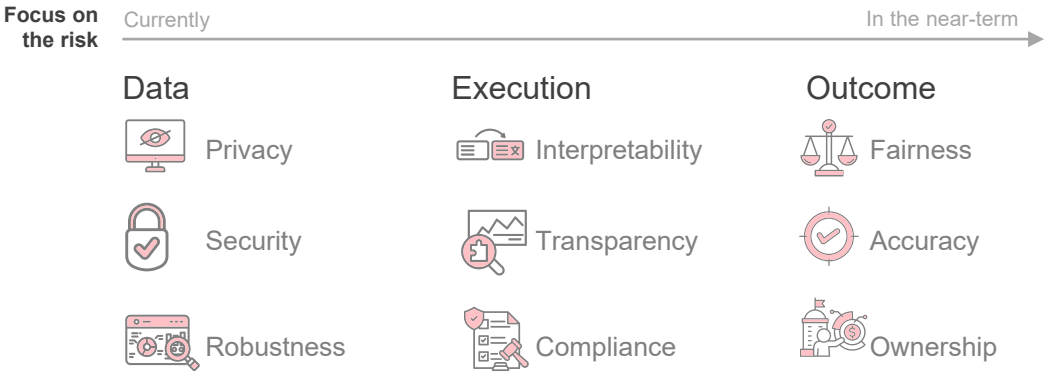

Exhibit 3 illustrates four risk categories the EU defined for generative AI applications and systems.

Exhibit 3: Risk categorization according to the EU AI regulations

Source: Everest Group (2024)

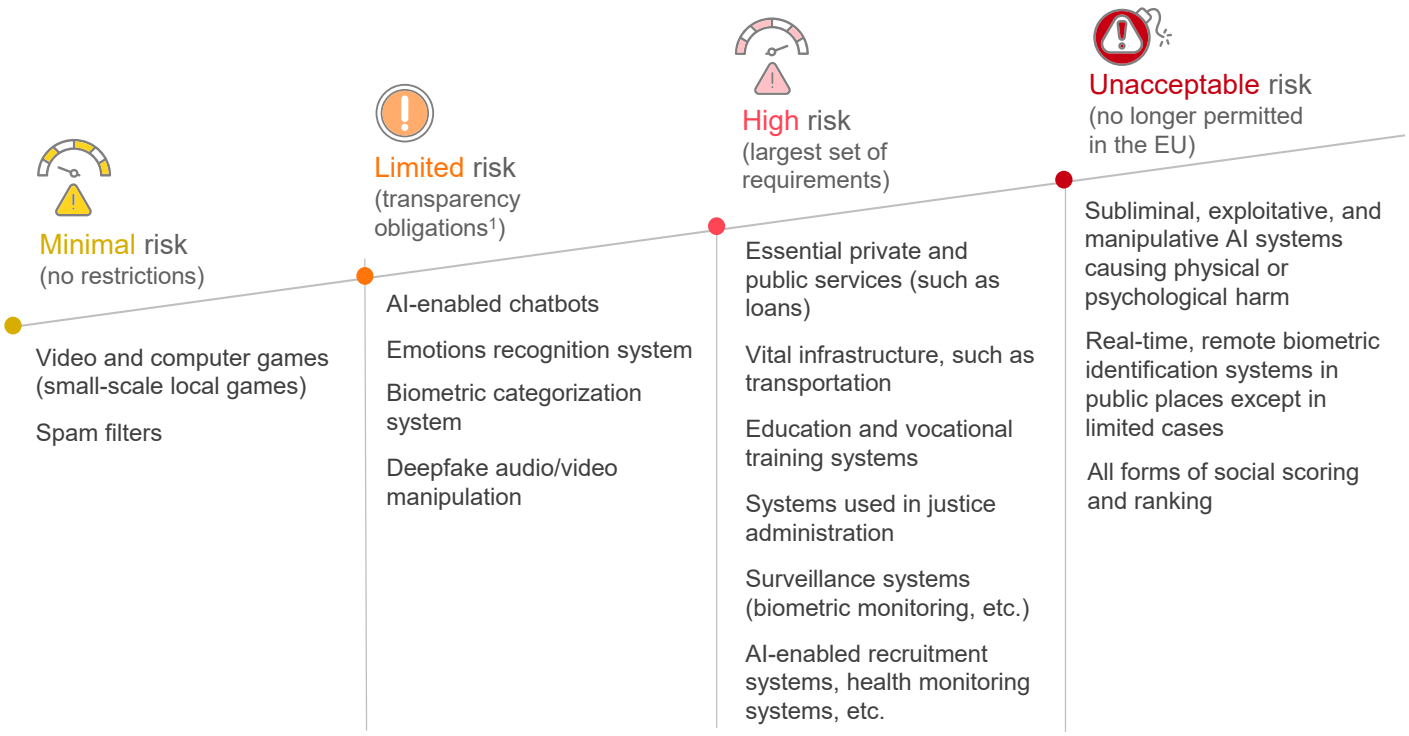

Note: Certain use cases may fall in the high-risk category as enterprises invest at scale

Enterprise perspectives on generative AI's risks

As we examine the types of use cases that enterprises are adopting, they are currently prioritizing ones with minimal to limited risk while steering clear of use cases that may expose them to higher risk levels.

85% of enterprises believe they are aware of generative AI's risks but are not prepared to handle them and, therefore, avoid investing in such use cases.

1 Transparency obligation: AI systems shall be designed so that users interacting with an AI system are informed that it is not a real human

This cautious approach minimizes risk but leads to lower returns. Enterprises limit their abilities to unlock generative AI's full potential by prioritizing use cases with minimal risk exposure. As a result, rewards from these minimal-risk use cases tend to be more incremental than transformative.

Exhibit 4 highlights the correlation between risk and rewards for generative AI use cases for Banking, Financial Services, and Insurance (BFSI) and Retail and Consumer Packaged Goods (RCPG) industries.

Exhibit 4: Generative AI use cases in the BFSI and RCPG industries

Source: Everest Group (2024)

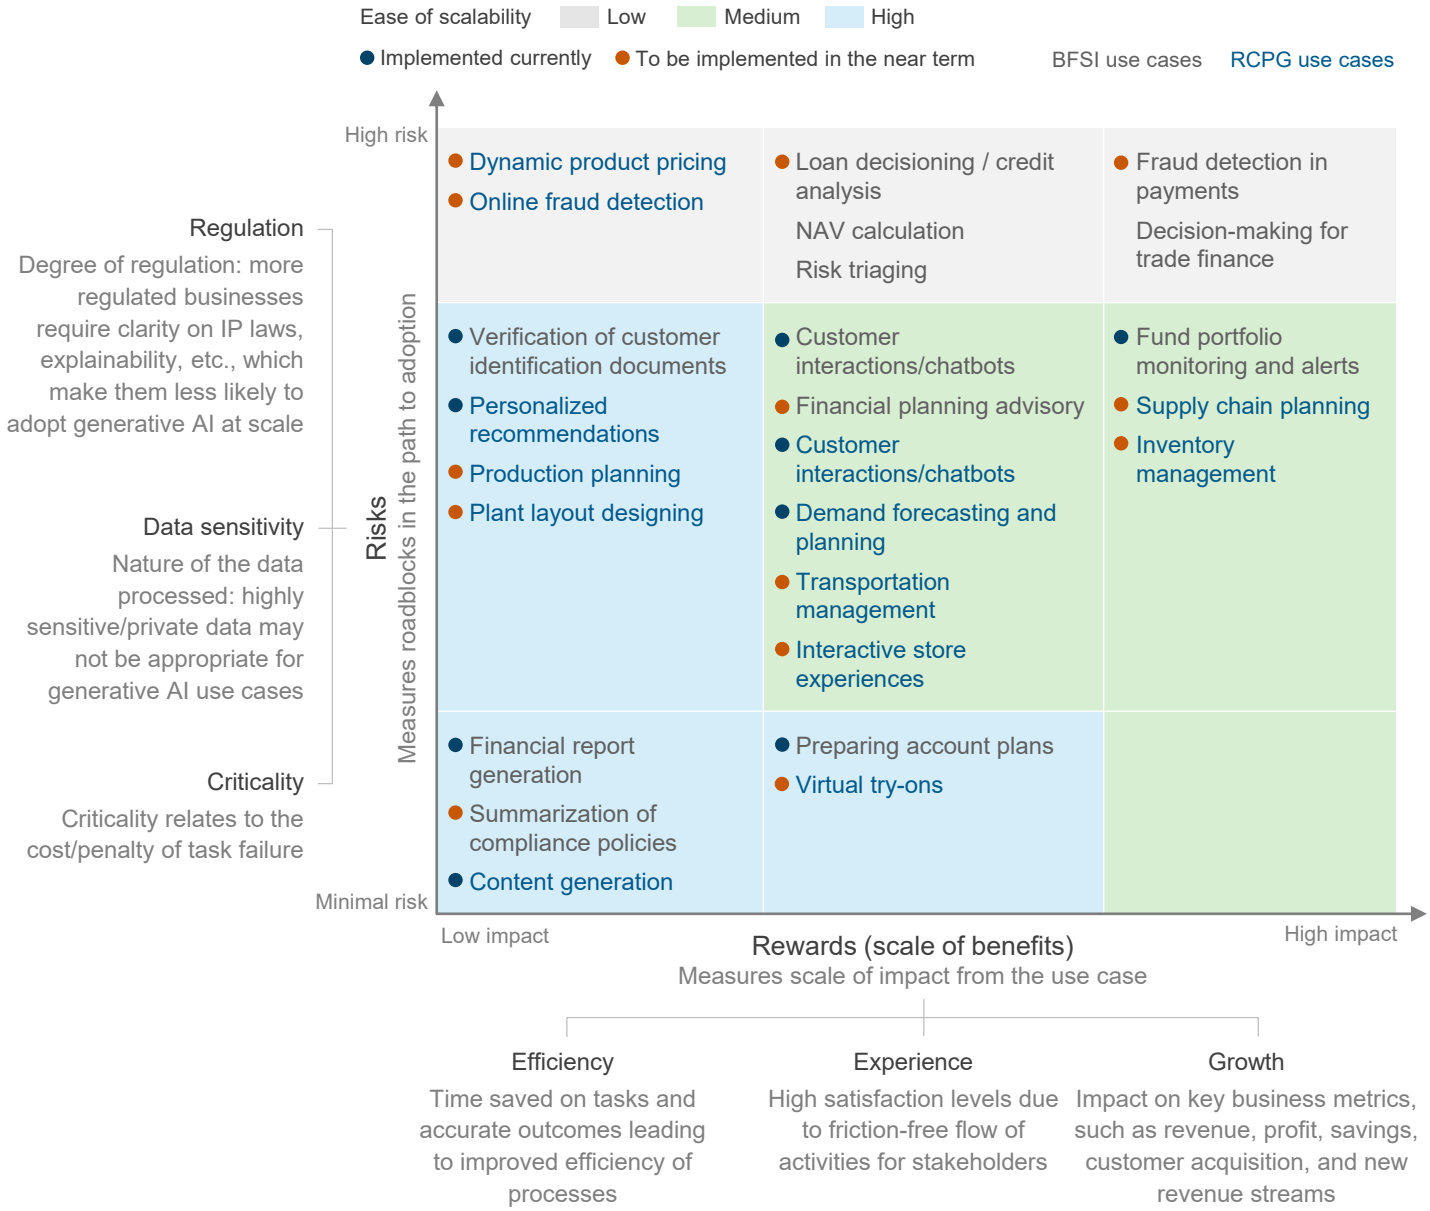

Exhibit 5 lists the correlation between risk and rewards for generative AI use cases for other industries.

Exhibit 5: Generative AI use cases across other industries

Source: Everest Group (2024)

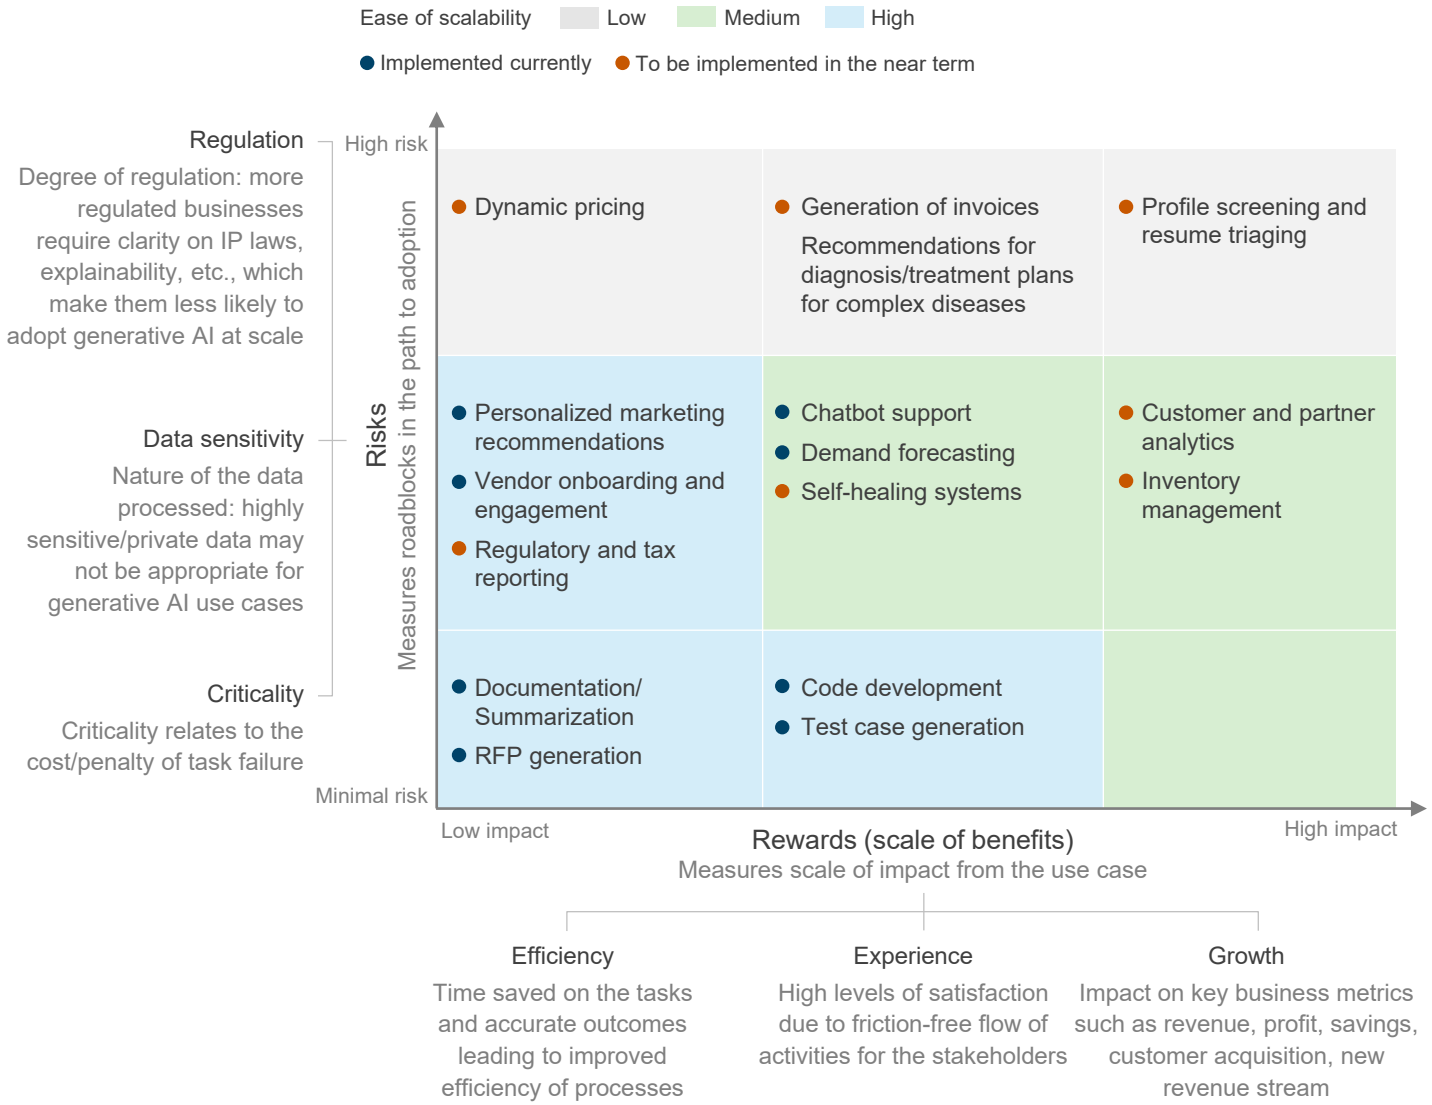

Enterprises must develop a robust risk management strategy to move beyond these limited rewards and fully capitalize on generative AI's current and future opportunities. They must embed responsible AI principles in their risk management strategy to ensure that the AI systems are developed and deployed ethically, safely, and transparently. This strategy should enable them to navigate higher-risk use cases more effectively, balancing risk and reward to unlock generative AI investments' full potential.

Understanding the need for risk management to maximize generative AI's rewards

Everest Group take

Enterprises must approach risk management comprehensively, addressing people, process, and technology dimensions. While several organizations have set up a risk management team and are making certain technology investments, more mature enterprises have targeted initiatives to prepare users and redesign risk management processes. Different teams' roles and responsibilities evolve as enterprises progress from planning to running operations powered by generative AI. A comprehensive approach across people, process, and technology ensures that enterprises are secure and well-prepared to navigate the risks associated with generative AI.

Most enterprises have invested in risk management by establishing dedicated teams, formulating enterprise-level data security guidelines, and partnering with trusted technology providers to support them in their generative AI initiatives. However, these efforts are often confined to the planning and implementation phases. The real challenge lies in the effective execution and compliance with risk management strategies as initiatives scale.

2 in 5 enterprises have advanced their risk management strategies and are seeing 25% higher rewards, including enhanced customer experience, increased productivity, and faster time to market.

Additional investments mature enterprises make for generative AI risk management

Mature enterprises typically invest in additional measures across people and process beyond technology investments, as depicted in Exhibit 6.

Exhibit 6: Enterprise investments in risk management strategies

Source: Everest Group (2024)

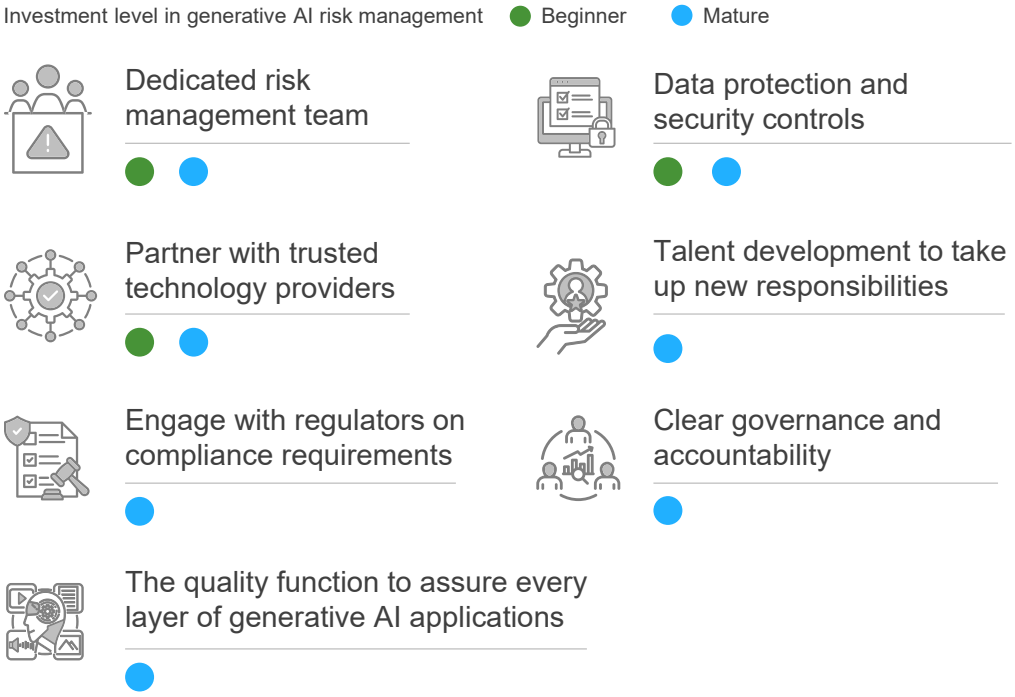

People

- Create additional roles, such as the Head of Responsible AI and the Chief Trust Officer (CTO), to ensure a 360-degree approach to risk management
- Invest in equipping users with the necessary skills to reduce adoption risks related to generative AI systems

Processes

- Engage and collaborate with AI regulation and compliance bodies to build a robust framework for risk management at the enterprise level
- Abide by relevant regulations, such as the European Union's Artificial Intelligence Act (EU AIA), the California Consumer Privacy Act (CCPA), the General Data Protection Regulation (GDPR), the Organization for Economic Co-operation and Development (OECD) Principles for Artificial Intelligence, the Global Alliance on Artificial Intelligence (GAAI), International Organization for Standardization (ISO) frameworks for AI, and the industry-level AI guidelines

- Involve the quality function to build and maintain generative AI systems' quality at every stage. The quality function is vital in these enterprises to address ethical concerns and mitigate bias in AI applications. By prioritizing data quality, model fairness, and transparency, enterprises can mitigate harmful outcomes and promote responsible AI development

These enterprise investments across people and processes have led to higher-order rewards, including better quality outcomes from generative AI applications, enhanced productivity, and improved customer experience.

100% of mature enterprises believe that the quality function will be significant in assuring generative AI's rewards.

Exploring the quality function's role in building generative AI ecosystems

> **Everest Group take**
>
> While the risk management and central AI teams will be pivotal in formulating and overseeing the execution of generative AI risk management strategies, the quality function will play a vital role in advocating these strategies throughout the generative AI use case implementation cycle. However, enterprises must upskill their quality functions in technical and business areas to build and maintain the quality of their generative AI applications.

A best-in-class approach to effective risk management necessitates collaboration between risk management, central AI, AI project, and quality teams. Each function plays a distinct and significant role in enhancing preparedness for managing risks and maximizing rewards from generative AI initiatives.

Exhibit 7 provides a detailed breakdown of different functions' roles across the plan, build, and run phases of the generative AI use cases implementation cycle.

Exhibit 7: Best-in-class approach for different functions to ensure effective generative AI risk management

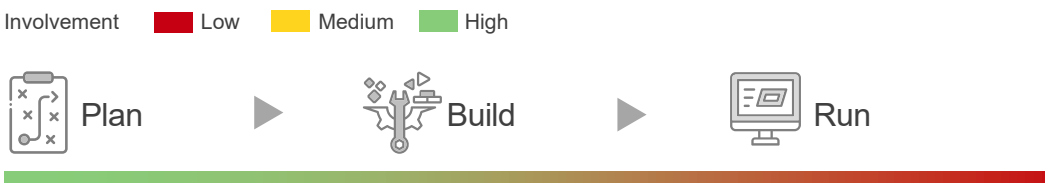

	Plan	Build	Run
The risk management function – the generative AI risk management strategy's architects	Identify potential risks for the organization and categorize risk levels Build a risk management guideline	Devise risk management strategies, such as data encryption, ethical guidelines for model development, and human oversight mechanisms Ensure compliance with relevant industry standards	Monitor the AI system performance to identify emerging risks or changes in existing systems Update risk assessment and management guidelines
The central AI function – the generative AI risk management strategy's guardians	Enable risk assessment and classification of generative AI use cases Evaluate and select appropriate tools for model development, training, and deployment	Oversee AI projects' execution Collaborate with development teams to ensure best practices are followed	Measure the effectiveness of generative AI solutions and identify improvement areas
The AI project function – the generative AI risk management strategy's executors	Select models that align with the use case and data characteristics Design the AI model's architecture Develop techniques to explain the model's decisions and reasoning	Manage data pipelines and ensure data is clean, accurate, and accessible for AI projects Assign labels or tags to data elements to offer context for AI training Develop, train, and deploy the AI model	Refine the model iteratively to enhance its performance
The quality function – the generative AI risk management strategy's advocates	Advocate risk management strategies across execution teams Embed quality by design principles	Ensure the AI system's proper integration with organizational systems Evaluate the model's ability to offer understandable explanations for its decisions Assure compliance with guidelines and policies for AI use Ensure training data quality and outcomes from the model	Model degradation testing Monitor for external attacks and security threats Detect potential biases or imbalances in the model's outcomes Assure compliance with guidelines and policies on an ongoing basis

> "Risk management due to generative AI adoption will be a collaborative effort of different enterprise teams. The quality function validates that there are no/minimal risks in the generative AI systems and applications."
>
> – IT leader of a large financial services firm

The risk management function: the generative AI risk management strategy's architects

- Formulate the risk management strategy and guardrails that other teams will follow
- Play a key role in the planning phase of the generative AI use case to identify risks before different generative AI use cases' implementation
- Coordinate with other teams to ensure risk management guidelines are followed and monitor the execution in the later part of the build and run phase

The central AI function: the generative AI risk management strategy's guardians

- Guide and ensure standardization across parameters, such as tools to be used, processes to be followed, and metrics to be tracked, in all the generative AI projects across the enterprise
- Involved throughout generative AI initiatives' implementation cycle as an observer to ensure adherence to risk management guidelines

The AI project function: the generative AI risk management strategy's executors

- High involvement in the plan and build stage to execute generative AI projects in line with the enterprises' generative AI risk management strategies and guidelines
- Coordinate closely with the central AI team to ensure that all best practices for execution are being followed

The quality function: the generative AI risk management strategy's advocates

- Limited but vital role during the planning phase of the generative AI use case but expands during the operational phase
- Early involvement in the planning phase to ensure that quality considerations are integrated into the design and promote risk management strategies in business units
- During the build phase, responsible for validating the data used to train the model, assessing the generative AI model's relevance, and ensuring the acceptability of the model's outcomes
- In the operational (run) phase, validate outcomes and continuously monitor the model's quality

Limiting the quality function's involvement to the stable state of a generative AI use case will hinder effective risk management during data preparation and model execution. This limitation results in lower-impact rewards, even after stabilizing the use case.

To maximize the potential rewards of generative AI applications, the quality function must promote a quality-first approach and instill a risk management mindset in the execution team from the planning phase.

The quality function's evolution with generative AI use cases' adoption across industries

Testing generative AI applications differs from traditional digital applications due to variable and evolving outcomes based on the model. These outcomes can change as the model learns and adapts, making it vital for quality teams to develop an approach to maintain outcome quality from generative AI applications and systems.

Exhibits 8 and 9 illustrate the quality function's evolved role in the banking and RCPG value chain to manage risks and assure generative AI's rewards.

Exhibit 8: The quality function's evolved role in generative AI-led use cases in the banking value chain

Source: Everest Group (2024)

Risk level ● Low ● Medium ● High

Process in the banking value chain	Customer servicing	Risk and regulatory compliance	Customer onboarding	Document extraction / summarization	Marketing / lead generation
Generative AI use cases	● Personalized business account insights ● Customer interactions/ chatbots ● Loan decisioning / credit analysis	● Synthetic data used for compliance testing and risk model development ● Fraud decisioning in payments ● Summarizing the bank's sustainability performance	● Loan contract generation (such as mortgages)	● Regulations/ Compliance/Policies summarization ● Extract information from contracts / legal documents	● Prepare account plans ● Track conversion rates and CSAT ● Compose personalized emails
Risks	Incorrect/incomplete training | Data bias in training data | Data security and privacy | Model hallucination | Ethicality of outcomes				
The quality function's evolved role	Chatbot testing Experience assurance Test for hallucination and explainability Outcome integrity	Compliance assurance Synthetic data generation	Error detection in generated documents	Data validation Knowledge injection	Data assurance Testing for bias
	Enable risk classification and risk evangelization, establish a Quality Management System (QMS), establish best practices to ensure traceability, continuous compliance monitoring				
Rewards	Improved customer satisfaction	Enhanced brand reputation	Improved customer experience	Enhanced efficiency	

Exhibit 9: The quality function's evolved role in generative AI-led use cases in the RCPG value chain

Source: Everest Group (2024)

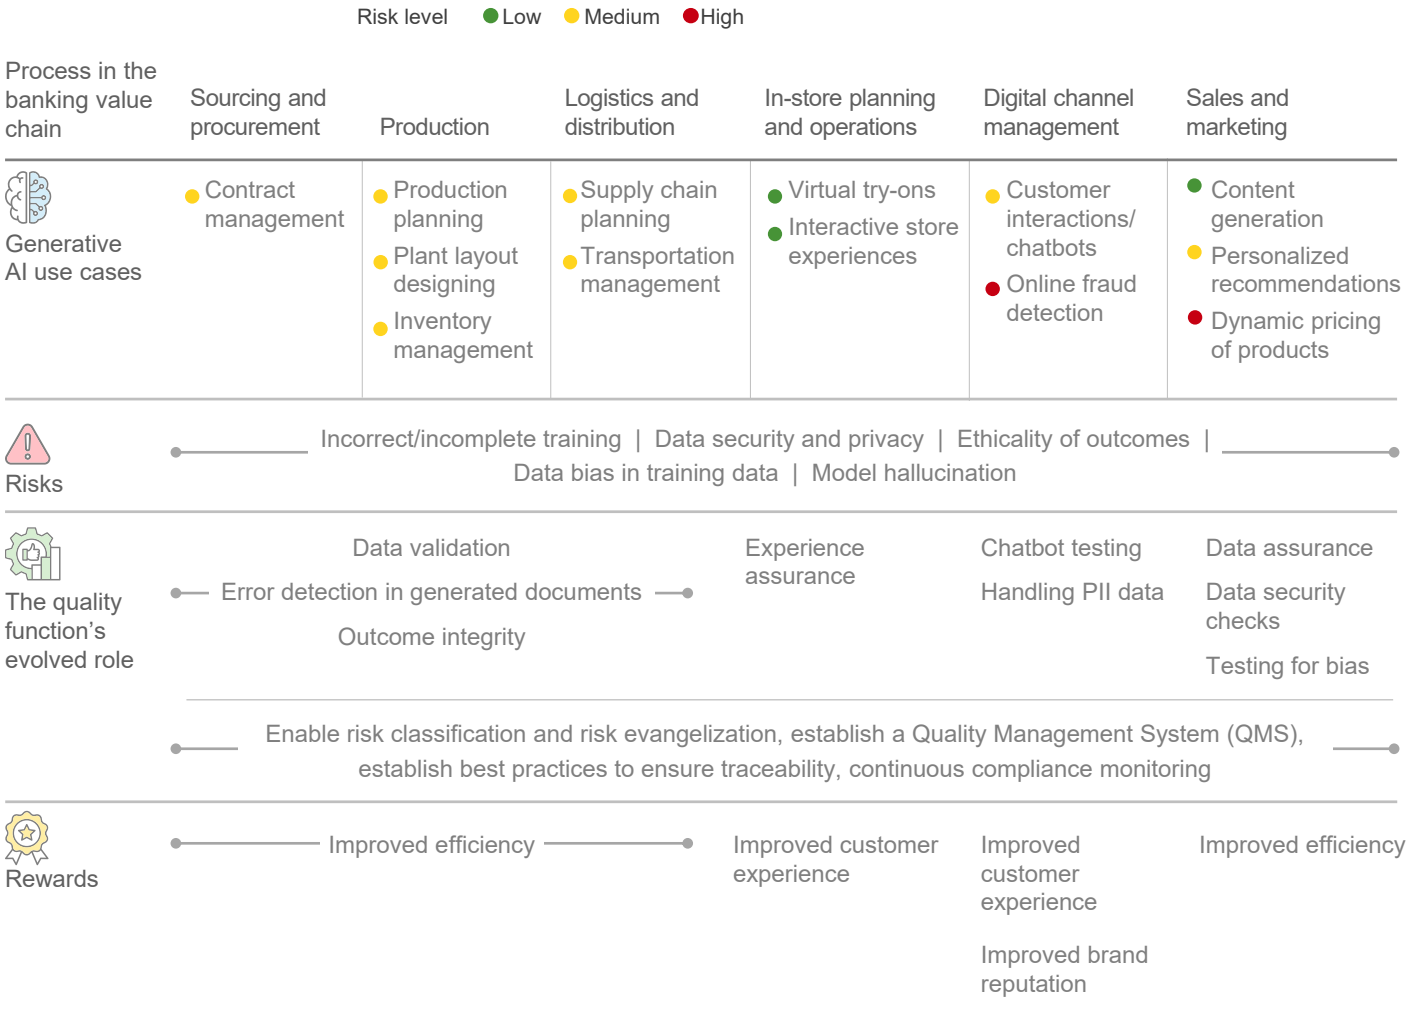

Though enterprise leaders believe that the quality function's role will evolve with the advent of generative AI, they also believe that the quality function is not ready to take up this added responsibility. Enterprises actively investing in generative AI use cases will also have to upskill the quality function to manage the risks of generative AI use cases and scale them.

80% of enterprises believe the quality function must enhance business and regulatory knowledge to ensure the benefits of generative AI use cases.

Enterprise roadmap for the quality function in a generative AI-driven business model

> **Everest Group take**
>
> Enterprises should adopt a **train-and-track** approach to address the gaps in the quality function to maintain the quality of generative AI applications, manage risks, and scale the generative AI initiatives. By boosting the quality function's capabilities and adjusting the performance metrics it tracks, enterprises can effectively manage risks and maximize rewards from their generative AI investments.

Enterprises should adopt a two-pronged **train-and-track** approach to close the capability gaps in the quality function, assure generative AI use cases and position them as custodians of AI rewards.

Exhibit 10 outlines the train-and-track approach for enterprises' immediate and near-term investments.

Exhibit 10: The train-and-track approach for enterprises' immediate and near-term investments

Source: Everest Group (2024)

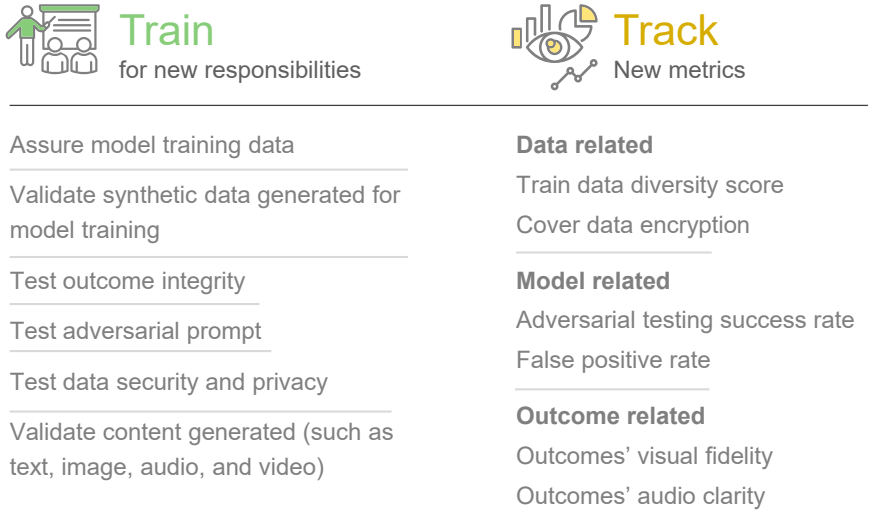

Train the quality function to deliver new responsibilities

Enterprises must upskill their quality function talent with AI-specific and business-oriented training to fulfill model and outcome validation responsibilities. A new skillset will be required to test generative AI systems and prevent potential quality issues effectively.

Exhibit 11 shows quality issues and the relevant skills enterprises must build to prevent them.

Exhibit 11: Quality issues and skills that enterprises need to build to prevent them

Source: Everest Group (2024)

Generative AI use case implementation cycle	Quality issues	Skills to be built
Plan 	Bias in training data	• Synthetic data generation and validation techniques • Data science in test • Data augmentation • Data preprocessing and cleaning
	Data confidentiality	
Build 	Model inaccuracies and drifting	• Model interpretability • Model regularization techniques • Model validation
	Block box nature of the model and decision-making	
	Hallucination	• Business knowledge • Ethical consideration knowledge
	Regulation and compliance	Knowledge of industry-specific compliance guidelines and regulations
Run 	Adversarial prompts	Adversarial defense techniques
	Outcome integrity	• Domain expertise • Explainability specialist

Track additional metrics for measuring the quality of generative AI applications and systems

Conventional testing metrics must be updated to track key parameters, such as model degradation, and detect potential biases or imbalances in generative AI outcomes. This testing ensures continuous monitoring of generative AI applications' evolving data, models, and outputs. The enterprise quality function must introduce new metrics across AI systems' layers to effectively monitor these parameters and realize the full benefits of generative AI.

Exhibit 12 depicts parameters enterprises must track to assure generative AI rewards for different areas.

Exhibit 12: Parameters that enterprises must track to assure generative AI rewards

Source: Everest Group (2024)

Generative AI use case implementation cycle	Area	Parameters
Plan 	Training data quality	Ensures the training data is: • Diverse • Complete • Correct • Bias free • Consistent
Build 	Model quality	• Accuracy measures the expected precision thresholds that the model should meet • Latency measures the time the model takes to respond to queries • Error rate tracks the percentage of incorrect or invalid responses the model generates
Run 	Model and data drifting	Continuous monitoring and refining the model with diverse inputs to identify improvement areas and fine-tune it accordingly
	Quality of content generated	• Coherence measures whether the generated content is logically consistent • Fluency assesses the AI-generated text's readability and naturalness • Relevance measures generated responses' relevance to the given prompts • Similarity compares the expected results with AI-generated outputs
	Risks in the content generated	Validates generative AI-generated content for: • Hateful and unfair content • Sexual content • Violent content • Self-harm-related content
	Overall application/ system performance	• Measures the generative AI application's/system's time to generate responses • Assesses the generative AI system's ability to complete specific tasks accurately and efficiently • Analyzes outputs generated in a specific timeframe to measure the throughput
	User satisfaction	• User feedback score collects qualitative and quantitative feedback from users regarding AI-generated content's relevance and quality • Net Promoter Score (NPS) measures user willingness to recommend the generative AI application to others

Other investments in the quality function

Apart from the train-and-track strategy, there are other best practices that the quality team should focus on to take up added responsibilities due to generative AI adoption seamlessly. These added responsibilities are:

- **Realign QA processes:** establish a Quality Management System (QMS) to ensure traceability and continuous risk monitoring across the generative AI applications and systems
- **Leverage AI for assurance of AI applications and systems:** identify and implement use cases where AI can be embedded for AI applications and systems assurance
- **Integrate generative AI tools with existing QA processes seamlessly:** oversee this integration to maximize AI's benefits while maintaining established testing protocols
- **Ensure human oversight in the development and deployment of generative AI applications and systems:** include this in the form of human-in-the-loop review, AI-assisted human decision-making, and human-AI collaboration platforms
- **Collaborate with the central AI, AI project, and development teams:** identify potential vulnerabilities in the generative AI applications and systems early in the development cycle

Conclusion

As enterprises strive to adopt generative AI in their technology landscape, it becomes essential for them to understand and acknowledge the risks while aiming for the rewards. To optimize generative AI's rewards, enterprises should implement generative AI use cases responsibly and prioritize investment in generative AI risk management strategies. These strategies will be a collaborative effort between multiple enterprise teams, such as risk management, central AI, AI project, and quality teams.

The quality function is well positioned to minimize risks and assure rewards of generative AI from the beginning of the generative AI use case implementation cycle. Enterprises that are mature in risk management investments and have embedded quality functions in every layer of their generative AI applications or systems are seeing more impactful rewards.

However, enterprises must upskill their quality functions to take on newer responsibilities of assuring generative AI applications or systems to get optimum rewards. They will also have to realign the performance metrics that the quality function used to track the quality and performance of the generative AI applications and systems.

In the future, **enterprises will increasingly need to expand their quality function's role as the custodian of assuring the rewards of generative AI to scale their generative AI initiatives seamlessly.**

Appendix

This section offers a glimpse into the characteristics of the 400 respondents that formed the basis of the study.

Exhibit 13 shows the survey respondents split by industry, revenue, and region.

Exhibit 13: Split of survey respondents by industry, revenue, and region (percentage of respondents)

Source: Everest Group (2024)

Split of survey respondents by industry

Banking	19%	Healthcare and life sciences	5%
Financial Services	19%	Manufacturing	4%
Insurance	19%	Energy and utilities	4%
Retail and consumer goods	16%	Others	8%
Software and technology	6%		

Split of survey respondent organizations by annual revenue

US$ 5-10 billion	41%
US$ 1-5 billion	36%
US$ 10-50 billion	13%
>US$ 50 billion	10%

Split of survey respondents by region

North America	45%
Europe	26%
United Kingdom	15%
Asia Pacific	4%
Rest of the world	10%

Everest Group®
With you on the journey

Everest Group is a leading research firm helping business leaders make confident decisions. We guide clients through today's market challenges and strengthen their strategies by applying contextualized problem-solving to their unique situations. This drives maximized operational and financial performance and transformative experiences. Our deep expertise and tenacious research focused on technology, business processes, and engineering through the lenses of talent, sustainability, and sourcing delivers precise and action-oriented guidance. Find further details and in-depth content at www.everestgrp.com.

This study was funded, in part, by Cognizant

For more information about Everest Group, please contact:

+1-214-451-3000

info@everestgrp.com

For more information about this topic please contact the author(s):

Chirajeet Sengupta, Managing Partner

chirajeet.sengupta@everestgrp.com

Alisha Mittal, Vice President

alisha.mittal@everestgrp.com

Swati Ganesh, Senior Analyst

swati.ganesh@everestgrp.com

Notice and Disclaimers