



HONG KONG MONETARY AUTHORITY
香港金融管理局

Generative Artificial Intelligence in the Financial Services Space

September 2024

Supported by



QUINLAN
& ASSOCIATES

Contents

Foreword	3
Introduction	5
1. Executive Summary	7
2. Evolvement of GenA.I.	9
3. GenA.I. Adoption and Impact in the Financial Services Industry	21
4. Regulatory Principles	38
5. Adoption Considerations and Suggested Practices	53
6. Facilitation Initiatives in Hong Kong	66
7. Way Forward	69
8. Appendix	71

Foreword

In recent years, the rapid evolution of technology has redefined the landscape of financial services, presenting new opportunities to enhance efficiency, decision-making, and customer experience. Among these technological advancements, Generative Artificial Intelligence (GenA.I.) stands out for its potential to reshape the way financial institutions operate, innovate, and engage with their customers. However, with this transformative power comes a critical responsibility: ensuring that the adoption of these technologies is both safe and ethical, while continuing to maintain the highest standards of trust and integrity.

The evolution of banking technology has been characterised by increasingly shorter innovation cycles. While the widespread adoption of automated teller machines (ATMs) took roughly 15-20 years since its introduction in the 1960s, subsequent banking innovations saw significantly accelerated adoption pace. Phone banking, introduced around the 1980s, took about 10 years to gain widespread acceptance, whereas internet banking, which emerged in the 1990s, achieved a more general level of adoption in approximately 5-8 years. With the unprecedented rapid uptake of GenA.I. applications such as ChatGPT by end-users, it is not surprising if there were an even more compressed innovation cycle for GenA.I.-powered financial services.

In response to the rapidly growing interest in Artificial Intelligence (A.I.) adoption from industry participants, the Hong Kong Monetary Authority (HKMA) published a series of whitepapers, including *Reshaping Banking with Artificial Intelligence*¹ (2019), *Artificial Intelligence in Banking*² (2020), and *Artificial Intelligence and Big Data in the Financial Industry*³ (2021), highlighting

the importance of A.I. for the financial services sector. These papers have contributed to the foundation of the HKMA's Fintech 2025 strategy to drive Fintech development in Hong Kong's financial services industry.

Fostering innovation within a well-defined risk management framework is key to the sustainable development of the financial services industry. In 2019, the HKMA issued a circular⁴ outlining high-level principles on risk management for the use of A.I. in banking. These principles, focusing on governance, application design and development, and ongoing monitoring and maintenance, remain highly relevant for GenA.I. applications today. While these foundational principles provide a solid and pertinent basis, the evolving nature of GenA.I. may require further practical guidance, which is being addressed through the development of the new Generative A.I. Sandbox rolled out jointly by the HKMA and the Cyberport.

This initiative is designed to provide banks with a risk-controlled environment in which they can experiment, test, and pilot novel use cases of GenA.I., while carefully considering the risk management and ethical implications of these advancements. Through these efforts, the HKMA aims to encourage innovation while promoting risk management, particularly in those areas highlighted in this paper in relation to data privacy, algorithmic bias, and cybersecurity. This aligns with the HKMA's broader commitment to encourage responsible A.I. adoption, which will enable financial institutions to unlock the full potential of GenA.I. while safeguarding the stability and resilience of the financial system.

¹ Hong Kong Monetary Authority. 2019. *Reshaping Banking with Artificial Intelligence*. (https://www.hkma.gov.hk/media/eng/doc/key-functions/financial-infrastructure/Whitepaper_on_AI.pdf).

² Hong Kong Monetary Authority. 2019. *Artificial Intelligence (AI) in Retail Banking*. ([https://www.hkma.gov.hk/media/eng/doc/key-functions/financial-infrastructure/Artificial_Intelligence_\(AI\)_in_Retail_Banking.pdf](https://www.hkma.gov.hk/media/eng/doc/key-functions/financial-infrastructure/Artificial_Intelligence_(AI)_in_Retail_Banking.pdf)).

³ Hong Kong Monetary Authority. 2021. *Report on "Artificial Intelligence and Big Data in the Financial Services Industry: A Regional Perspective and Strategies for Talent Development"*. (<https://www.hkma.gov.hk/eng/news-and-media/press-releases/2021/10/20211028-3/>).

⁴ Hong Kong Monetary Authority. 2019. *High-level Principles on Artificial Intelligence*. (<https://www.hkma.gov.hk/media/eng/doc/key-information/guidelines-and-circular/2019/20191101e1.pdf>).

This paper seeks to provide a comprehensive overview of the role GenA.I. can play in the financial sector, as well as the potential use cases and challenges that institutions should consider when integrating these technologies into their operations. It also details diverse approaches to enabling the adoption of GenA.I., focusing on practical guidance, industry collaboration, and capacity building.

As we look ahead, the HKMA remains committed to creating an environment where financial institutions can leverage the power of GenA.I. responsibly, while ensuring that Hong Kong continues to be a leader in Fintech innovation. The HKMA believes that by working together, we can unlock the full potential of GenA.I., while upholding the integrity, resilience, and stability of our financial system.

We would like to extend our heartfelt thanks to the Insurance Authority (IA), the Mandatory Provident Fund Schemes Authority (MPFA), financial institutions, industry stakeholders, and participants who contributed their insights to this paper. Collaboration is critical as we embark on this journey to shape the future of financial services through GenA.I.

“



Fostering innovation within a well-defined risk management framework is key to the sustainable development of the financial services industry.

Mr Arthur Yuen, Deputy Chief Executive of the HKMA

”

Introduction

The financial services sector is entering a new era of transformation, driven by the rapid advancement and integration of emerging technologies. To ensure Hong Kong remains at the forefront of global Fintech innovation, the HKMA has identified five strategic focus areas: Wealthtech, Insurtech, Greentech, A.I., and Distributed Ledger Technology (DLT). These focus areas are pivotal to the HKMA's Fintech 2025 strategy and its underlying "All banks go Fintech" initiative, aimed to future-proof the financial services industry and ensure Hong Kong remains a global leader in Fintech innovation.

In its ongoing mission to promote responsible adoption of technology, the HKMA has recently launched the FiNETech series together with the other local financial regulators. This initiative fosters cross-sectoral Fintech collaboration and provides a platform for financial institutions to engage with cutting-edge technologies. The latest edition, FiNETech2, of the series was held in August 2024, focusing on exploring A.I., including GenA.I., for its ability to revolutionise key aspects of banking and financial services.



FiNETech2 - Into the A.I. verse

As financial institutions are entering the next phase of A.I., including the exploration of GenA.I. adoption, the HKMA recognises the importance of an interactive and iterative approach. In response, the GenA.I. Sandbox has been introduced to provide a risk-controlled environment where institutions can develop, test, and pilot innovative GenA.I.-based solutions in real-world banking scenarios. Through the Sandbox, participants will receive early supervisory feedback and guidance from the HKMA, helping them to innovate GenA.I. responsibly. The Sandbox trials will focus on enhancing three priority areas that are expected to benefit most from GenA.I. risk management, anti-fraud measures, and customer experience.

With the first cohort of the GenA.I. Sandbox⁵ now open for applications, the HKMA looks forward to collaborating closely with banks and relevant stakeholders. Leveraging the Sandbox, the ecosystem can jointly explore novel use cases that drive responsible innovation, improve operational efficiency, and enhance the effectiveness of financial systems. The HKMA also reiterates its commitment to supporting such collaborative partnerships between financial institutions and Fintech firms, drawing on the expertise of both sectors to unlock the full potential of GenA.I. technologies.

This research paper is another attempt to provide a building block to help banks and other financial institutions to sail through the A.I. journey by sharing insights and practical experiences. It offers an indicative roadmap for integrating GenA.I. solutions into financial operations. The HKMA recognises that the adoption of A.I. technologies demands not only innovation but also a robust governance framework. Accordingly, the HKMA advocates for a comprehensive approach that includes risk management, system integration, and continuous monitoring to ensure the safe and ethical deployment of these technologies.

⁵ Hong Kong Monetary Authority. 2024. *Generative Artificial Intelligence Sandbox*. (<https://www.hkma.gov.hk/media/eng/doc/key-information/guidelines-and-circular/2024/20240920e1.pdf>).

Moreover, the HKMA is placing a strong emphasis on capacity building to ensure that the financial services workforce is equipped to navigate the complexities of GenA.I. adoption. Through training initiatives and talent development programmes, the HKMA aims to build a future-ready workforce with the skills necessary to embrace A.I. technologies responsibly and adaptively.

As we move forward, the HKMA remains committed to fostering a Fintech ecosystem that balances innovation with supervisory oversight, ensuring that the benefits of GenA.I. can be realised without compromising the stability, security, or ethical standards of the financial system. By working closely with fellow regulators, industry stakeholders, financial institutions, and technology providers, Hong Kong is well-positioned to lead the way in the responsible integration of A.I. and GenA.I. with the global financial systems to achieve effective outcomes.

“

As financial institutions are entering the next phase of A.I., including the exploration of GenA.I. adoption, the HKMA recognises the importance of an interactive and iterative approach.

Ms Carmen Chu, Executive Director (Banking Supervision) of the HKMA

”

1. Executive Summary

The emergence of GenA.I., powered by large language models (LLMs) that facilitate seamless human-A.I. interaction, marks a pivotal advancement in the evolution of A.I. technology. At its core, GenA.I. possesses the capabilities to generate human-like, novel outputs across a variety of media such as text, images, audio and video with profound implications for numerous industries. Among these, the financial services sector stands to be particularly transformed by GenA.I., given its potential to revolutionise operational efficiency, risk management, and customer engagement.

GenA.I.'s rapid growth is fuelled by advances in computational power and increasingly sophisticated model architectures, enabling financial institutions to streamline key functions. Two key areas are knowledge management and creative content management processes, which are enhancing operational efficiency.

To gauge the current adoption efforts of GenA.I., the HKMA conducted an industry-wide survey participated by 137 industry practitioners and in-depth interviews with 16 organisations across the technology, banking, securities, and insurance sectors. Findings reveal that most financial institutions are currently focusing on internal applications that drive employee productivity. These include innovations aimed at automating routine tasks, improving workflow efficiencies, and enhancing knowledge sharing. The adoption of GenA.I. for customer-facing applications remains nascent. Concerns around regulatory compliance, operational deployment, and the resource-intensive nature of GenA.I. solutions are cited as key barriers to more widespread adoption.

In response to rapid advancements in GenA.I., regulators globally have been revisiting or updating their frameworks to ensure that the deployment of A.I. technologies is both responsible and ethical. Some jurisdictions, including the European Union (EU), Hong Kong Special Administrative Region (SAR), Mainland China, Singapore, United Kingdom (U.K.), and United States (US) are refining regulations or supervisory guidance to strike a balance between innovation and consumer protection. These frameworks emphasise critical principles such as:



Governance & Accountability



Fairness



Data Privacy & Protection



Transparency & Disclosure

While the potential of GenA.I. to transform the financial services sector is immense, its adoption also introduces several key risks that financial institutions must address to ensure safe and responsible deployment. These include concerns around inaccuracy or hallucination, where A.I. models may generate incorrect or misleading information, as well as data privacy risks due to the large datasets often required by GenA.I. systems. Additionally, risks related to algorithmic bias, cybersecurity vulnerabilities, and operational challenges such as system integration and scalability must be carefully managed to avoid negative impacts on both operations and reputation.

Against this regulatory backdrop, financial institutions are encouraged to establish a robust GenA.I. governance structure that ensures compliance and facilitates the responsible deployment of new solutions. This governance framework should include the development of business cases and proof-of-concepts that are complemented by comprehensive strategies for risk management, system integration, and continuous monitoring. A well-designed feedback loop across all phases of adoption is essential to drive ongoing improvements and maximise the value of GenA.I. initiatives while mitigating risks.

Recognising the resource intensity required to implement A.I and GenA.I. solutions, Hong Kong has introduced a range of facilitation initiatives to support financial institutions. These include access to high-performance computing resources, financial subsidies, talent development programmes, and promotional events aimed at fostering innovation and building a skilled workforce capable of leveraging GenA.I. technologies.

2. Evolvment of GenA.I.


A.I. has experienced significant advancements in recent years, driven by the exponential growth in data generated by our increasingly digitalised world while being supported by ever-improving advancements in computing hardware.

Historically, A.I. software, while a powerful tool, was designed to handle specific tasks with more limited and less intuitive user interactions. However, the advent of LLMs – such as GPT by OpenAI, Llama by Meta, and Gemini by Google – has broadened access, enabling individuals, irrespective of technical expertise, to interact with A.I. systems seamlessly. This development marks a new era of global A.I. adoption.

LLMs have enabled various A.I. models to not only automate tasks but also generate original outputs, such as text, code, images, audio, and video, domains once considered exclusive to human creativity. These A.I.-generated outputs now frequently meet or even surpass human standards. The introduction of ChatGPT in 2022 marked a significant turning point in the A.I. revolution, igniting widespread public interest and enthusiasm for GenA.I. solutions.

According to a report by Bloomberg Intelligence⁶, the market for GenA.I. is projected to grow at a compound annual growth rate of 42% over the next decade, underscoring the profound interest in and transformative potential of these technologies (see Figure 1).

Figure 1: GenA.I. Market Potential



Source: Bloomberg Intelligence

⁶ Bloomberg. 2024. *Generative AI to Become a \$1.3 Trillion Market by 2032, Research Finds*. (<https://www.bloomberg.com/company/press/generative-ai-to-become-a-1-3-trillion-market-by-2032-research-finds/>).


The financial services industry is actively exploring GenA.I. adoption opportunities, such as personalising customer interactions, optimising financial products, and enhancing decision-making processes.

2.1. A.I. Supporting Pillars and GenA.I. Implications

Before exploring the concept of GenA.I., it is important to recognise that GenA.I. is a subset of A.I., meaning its functionalities are based on the same foundational principles as A.I. technology. Therefore, having a solid understanding of how A.I. operates is key to fully grasping the concept of GenA.I. The functionality of A.I. is built upon and supported by four key pillars that govern the training and operation of its models (see Figure 2).

- 1 Training Data:** The dataset used to train and refine A.I. models, providing the foundational content and relevant context required for generating outputs.
- 2 Training Methods:** The algorithms and techniques employed to optimise A.I. models, ensuring they learn effectively from the data and improve their performance over time.
- 3 Model Architecture:** The design configurations and structural frameworks of A.I. models, which dictate how they process information and generate output.
- 4 Computational Power:** The processing and memory capabilities required to train and run A.I. models, enabling them to handle complex tasks and large datasets efficiently.

Figure 2: A.I. Supporting Pillars



Source: Subject matter expert interviews, Quinlan & Associates analysis


Training Data & Method

Before A.I. solutions can be developed and commercialised, the solution provider begins by training the model using a vast dataset, commonly referred to as corpus data, which serves as the foundational 'content' for the model. Over time, the output is further refined through exposure to additional data from end-users, external sources, and ongoing evaluations, providing additional 'context' for the model's output generation. This continuous refinement allows the model to generate outputs

that more closely align with desired benchmarks. However, it is important to note that while the model benefits from this influx of new data, the original corpus and the foundational architecture of the model typically remain unchanged.

The continuous improvement process of A.I. models is achieved through three primary training methods: (1) supervised learning, (2) unsupervised learning, and (3) reinforcement learning. Each method differs in how it interacts with the training data and adapts to new inputs (see Figure 3).

Figure 3: A.I. Training Data & Methods



Source: MIT Sloan School of Management, Quinlan & Associates analysis

Supervised learning involves the provision of labelled training data to the model, where the correct output is already known, enabling the model to learn the mapping function between inputs and outputs. For instance, in the financial services domain, supervised learning algorithms are extensively used in the detection and prevention of fraudulent activities. The training data would consist of examples of both fraudulent and non-fraudulent wire transfers, each explicitly labelled with its respective category. Subsequently, the model is trained to extrapolate this knowledge to new, unlabelled wire-transfer data, effectively classifying it based on the learnt patterns. The breadth of potential applications for supervised learning is substantial, extending to tasks such as sentiment analysis, wherein textual data is

systematically classified according to its emotional tone, thereby providing insights into consumer sentiment.

Unsupervised learning operates in the absence of labelled data, allowing the model to autonomously explore and identify latent structures within the dataset without predefined outputs. This method is particularly advantageous for exploratory data analysis, clustering, and anomaly detection. For example, a model trained with unsupervised learning might cluster consumers based on their raw purchasing behaviour data, without prior knowledge or labelling of existing segments. Such clustering can be instrumental in other potential applications such as customer profiling, enabling financial institutions

to develop tailored credit scores for distinct consumer segments, and recommendation systems that suggest products based on inferred user behaviour patterns.

Reinforcement learning diverges from the aforementioned methods by using feedback from its environment to iteratively refine the model's behaviour. This approach is particularly effective in controlled settings, such as financial markets, where models are trained to perform specific tasks by learning from the rewards and penalties received by their actions. Potential use cases for reinforcement learning include autonomous trading systems, which navigate the complexities of the financial market by optimising the balance of risk and return.

Collectively, these learning methodologies are fundamental to the continuous improvement of A.I. systems, enabling them to assimilate new data and experiences in increasingly sophisticated and nuanced ways.

One of the key advancements in A.I. is the introduction of retrieval-augmented generation (RAG). While not directly applicable to model training, this approach enables A.I. models to access data beyond their original training corpus, drawing from both internal sources (e.g., user input) and external sources (e.g., the web). This allows models to generate more accurate, relevant, and up-to-date responses without altering their fundamental structure. RAG enhances a model's ability to meet user needs more effectively, making it adaptable and responsive to real-time information while preserving its core model.

A.I. Model Architecture


All deep learning models that underpin A.I. are structured around artificial neural networks comprising layers of interconnected nodes. This architecture enables the model to learn complex relationships between data points through numerous iterations, thereby becoming increasingly sophisticated over time.

These models can be conceptualised as complex mathematical functions that execute computations on input data to generate an output. At their core, these models consist of nodes and parameters, which are critical in the information processing pipeline. Nodes represent individual computation units within the model, while parameters, specifically weights and biases, determine how input signals are transformed as they propagate through the network. Weights regulate the strength of the connections between nodes, influencing the extent to which one node impacts another. Biases, on the other hand, facilitate adjustments in the output independent of the input, providing the model with additional flexibility.

The relationship between input and output is mediated through a feed-forward mechanism. In this process, input data traverses multiple layers of calculations, with each layer applying distinct weights to the inputs. This hierarchical approach allows the model to cultivate a deeper and more nuanced understanding of the data, with hidden layers capturing complex, non-linear patterns and interactions within the input. Some of the modern neural network architectures, however, deviate from the purely feed-forward structure, particularly those with memory or attention mechanisms (e.g., recurrent networks or transformers). These models introduce additional dynamics, such as self-attention, which enhance the model's ability to process and retain information.

A critical aspect of training an A.I. model is the back-propagation algorithm. This technique initiates by calculating the error through a comparison of the model's output to a predefined benchmark or expected result. Upon identifying the error, the model adjusts its parameters with the objective of minimising this discrepancy. The recalibration process is computationally intensive, often involving billions of parameters, each necessitating precise calculations to determine the optimal adjustments. This iterative refinement is indispensable for enhancing the model's accuracy and overall performance, allowing it to learn more effectively from the data it processes (see Figure 4).

Figure 4: A.I. Model Architecture



Source: Google, Quinlan & Associates analysis

Processing Hardware

The final crucial element in the building of A.I. systems is computational power, which has experienced rapid and transformative growth over the past decade,

driven by significant advancements and specialisation in relevant hardware components (see Figure 5).


Figure 5: Hardware Development

Processing Hardware Comparison

			
	Central Processing Units	Graphic Processing Units	Tensor Processing Units
Description	Execute main instructions and perform general-purpose computations	Handle highly parallel workloads with many smaller, highly efficient processing cores	Specialise in processing tensor operations, which are fundamental to deep learning models
Introduction	1971 Intel 4004 by Intel	1999 Geforce256 by NVIDIA	2015 Tensor Processing Units by Google
Specialisation	Sequential Processing	Simultaneous Processing	
Functions	- Main system operations - Application processing	- Calculation in parallel - Computation-intensive tasks	- Large-volume processing - Execution of neural networks
Examples	- Intel Core i9 (Intel) - Ryzen Threadripper 7980X (AMD)	- RX 590 (AMD) - Tesla V100 GPU (NVIDIA)	- Tensor Processing Units (Google) - Tensor Core (NVIDIA)
	Ensure an A.I. system's overall management and control	Applied to A.I. training for parallel processing abilities	Speed up matrix calculations for A.I. training by design

Single-Chip Inference Performance

Int 8 TOPS* for NVIDIA chips, 2017-2023



*A performance metric for A.I. hardware that measures the speed of processing operations (such as arithmetic computations) in neural networks, particularly when using 8-bit integer precision

Source: Amazon, Google, Intel, NVIDIA, Quinlan & Associates analysis

Central processing units (CPUs) have served as the primary computational engines in computers since their inception in 1971. CPUs are responsible for the overall management and coordination of computer systems, executing instructions sequentially and performing a wide range of complex calculations. However, the inherently sequential processing capabilities of traditional CPUs are often inadequate for the compute-intensive tasks required for training and deploying modern A.I. models.

While CPUs excel at general-purpose tasks, the demands of A.I. necessitate extensive parallel processing, a domain in which CPUs are not inherently designed to operate efficiently. This limitation has catalysed the development of specialised hardware accelerators, such as graphics processing units (GPUs), which are better suited to the simultaneous processing requirements inherent in A.I. workloads.


Introduced by NVIDIA in 1999, GPUs were initially designed to accelerate computer graphics and image processing by leveraging numerous smaller, highly efficient processing cores. It was subsequently discovered that the parallel processing capabilities of GPUs were ideally suited to the computational

demands of training A.I. models, which involve processing vast volumes of data concurrently. This realisation significantly accelerated the development and sophistication of modern A.I. systems.

Beyond CPUs and GPUs, A.I.-specific accelerators have been developed to optimise the complex computations involved in training A.I. models. A prominent example is Google's Tensor Processing Units (TPUs), introduced in 2015. TPUs feature a custom hardware architecture specifically designed to accelerate low-level operations on multi-dimensional data, which are often the most computationally intensive and time-consuming parts of A.I. systems. These A.I.-specific accelerators achieve significantly higher performance and energy efficiency for A.I. workloads compared to general-purpose CPUs and GPUs.

Since the first transformer model was introduced by Google⁷ researchers in 2017, the number of parameters, which is essential for capturing intricate patterns in data, has grown exponentially. For example, the average number of parameters has increased from 373 million in 2017 to 139 billion in 2024, representing a 373-fold growth (see Figure 6).

Figure 6: Average Number of Parameters in Notable A.I. Systems



Source: Epoch, Quinlan & Associates analysis

⁷ Google. 2017. *Attention Is All You Need*. (<https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>).

As the parameter count rises, so does the model's capacity, along with the computational power required to train and operate it. The introduction of specialised processing units, such as NVIDIA's H100 series⁸, which has demonstrated a 32-fold improvement in inference tasks (i.e., making predictions or decisions based on new data) compared to its predecessor from 2017, has provided the computational power necessary for increasingly sophisticated A.I. models.

2.2. Key Concepts of GenA.I.

GenA.I., as the name suggests, is a transformative category of A.I. that excels at 'generating' novel outputs through the utilisation of generative models. These models are designed to discern patterns and structures from extensive training data.


At its core, GenA.I. leverages advanced algorithms within these generative models to assimilate knowledge from vast datasets. Unlike traditional models, which may produce outputs closely mirroring their input data, generative models have the capability

to create entirely new outputs. This enables them to produce more coherent and contextually relevant outputs in response to prompts, effectively mimicking human-like understanding and creativity.

LLMs, a type of generative model, can produce contextually relevant language and code outputs, such as summarising long documents or suggesting suitable source code for software development, with detailed contextual reasoning communicated in a human-like manner. These capabilities are widely applicable for enhancing everyday productivity across various industries and departments.

Beyond language and code generation through LLMs, there are diffusion models that extend to various creative domains, including images, audio, and video. Within seconds, these generative models can effortlessly create and produce compelling media, such as dynamic characters or avatars, audio snippets, and visually engaging video content. These capabilities are particularly useful for initiatives like press releases and marketing efforts by institutions (see Figure 7).

Figure 7: Generative Artificial Intelligence



Source: OpenAI, Quinlan & Associates analysis

⁸ NVIDIA. 2024. *Wide Horizons: NVIDIA Keynote Points Way to Further AI Advances*. (<https://blogs.nvidia.com/blog/hot-chips-daily-research/>).

GenA.I. is a collective term that encompasses various types of generative models, such as LLMs and diffusion models, capable of creating outputs ranging from language to video. However, it is important to note that these generative models, which support different domains, are distinct technologies that often complement one another, with LLMs and diffusion models running together to support integrated multimodal GenA.I. systems. Understanding this distinction and relationship is crucial as the financial industry, along with other sectors, explores the diverse use cases of this technology.

Large Language Models

LLMs have significantly advanced natural language processing by capturing the semantic meaning of words and generating contextually relevant outputs. This advancement is facilitated by the seamless interaction of several key components within LLMs, including word embedding, contextual encoding, and output generation (see Figure 8).

Figure 8: Large Language Model Mechanism

LLM MECHANISM	SIGNIFICANCE
1 WORD EMBEDDING Capture a word's semantic meaning with vectors	<i>Similar to previous models / architectures</i>
2 POSITION ENCODING Encode the words' positions in the vectors	<i>Factors in the word's sequence in the prompt to enable simultaneous vector processing for rapid response</i>
3 ATTENTION BLOCKS Encode weighted relationships among words*	<i>Identify the meaning of each word in the context to deeply comprehend the intent behind the prompt</i>
4 VECTOR CALCULATION Calculate vectors representing output elements	<i>Similar to previous models / architectures</i>
5 OUTPUT CONVERSION Convert vectors back to a suitable type of output	<i>Similar to previous models / architectures</i>



LLM allows accurate and seamless communication between humans and computers

*The model builds contextual representations of the meaning of a word (represented by tokens) by attending to every single word in the context and maps relationships between tokens within the same sequence (i.e., self-attention) and across tokens from other sequences (i.e., cross-attention) to integrate information from surrounding tokens.

Source: Google, Quinlan & Associates analysis

LLMs initiate the process by mapping words to high-dimensional vectors through word embedding. This technique encodes relationships and similarities between words, forming the foundation for semantic understanding. Building on this, LLMs enhance coherence and contextual relevance by encoding the order and relative positions of words and dynamically assigning weights to them. This weighting system allows the model to emphasise or de-emphasise words based on their significance within a given context.

At the core of LLM architecture are transformers, a fundamental innovation that has revolutionised language processing. Unlike traditional neural networks, which process data sequentially, transformers utilise a mechanism called “self-

attention.” This mechanism enables the model to evaluate the importance of each word in a sentence relative to every other word simultaneously. By processing information in parallel, transformers can capture complex dependencies and relationships within the text, regardless of the distance between words.

As such, transformer models have become the leading architecture for LLMs due to their versatility and effectiveness in handling various text-based tasks. Three primary types of transformer models have been developed to cater to different LLMs: (1) encoder-only models, (2) decoder-only models, and (3) encoder-decoder models (see Figure 9).

 **Figure 9: Types of Transformer Models for GenA.I. Solutions**

Transformer	Understandability	Text Classification	Text Generation	Text Summarisation	Text Translation	Notable Examples
Encoder-Only	Input only	✓	✗	-	✗	• Google (BERT, RoBERTa, DistilBERT) • AWS Bedrock (Titan)
Decoder-Only	Input & Output	✗	✓	-	✗	• OpenAI (GPT Series) • Meta (LLaMA) • DeepMind (Gopher)
Encoder-Decoder	Input & Output	-	-	✓	✓	• Google (T5) • Meta (BART, mBART)

✓ High Applicability - Moderate Applicability ✗ Limited Applicability

Source: Stanford University, Quinlan & Associates analysis

- **Encoder-only Models:** These models are designed primarily for understanding and representing text. They focus on encoding the input text into a fixed-length representation that captures the contextual information, making them highly effective for tasks that require text classification and comprehension. Notable examples include Google’s BERT and AWS’ Titan, powered by Bedrock.
- **Decoder-only Models:** These models are designed for text generation tasks. They predict the next word in a sequence based on the previous words, making them unidirectional but highly effective for generating coherent and contextually appropriate text. Notable examples include OpenAI’s GPT series, Meta’s LLaMA, and DeepMind’s Gopher.
- **Encoder-decoder Models:** These models are designed to handle tasks that require both understanding and generating text. The encoder processes the input text to understand its output and structure, while the decoder generates the output text based on the encoded representation. This architecture is well-suited for tasks that involve transforming text from one format to another, such as translation or summarisation. Notable examples include Google’s T5 and Meta’s BART.

Each type of model has distinct characteristics and is optimised for specific applications, ranging from text classification to translation. Understanding the unique strengths and limitations of these models is important when leveraging their capabilities in diverse real-world scenarios. These models can also be

further expanded by supporting multiple languages, known as multilingual models, and by adopting sparse models, which optimise computational costs for efficient large-scale tasks.


The Role of LLMs in the overall GenA.I. Landscape

The ability to understand and manage long-range dependencies enables transformers to produce coherent, contextually relevant outputs. As a result, transformers are central to the robust performance of modern LLMs, enabling them to excel in a diverse array of natural language processing tasks with remarkable accuracy and fluency. Once semantic understanding is established, LLMs perform complex computations to combine information from various

words and contextual cues, ultimately transforming these into meaningful and contextually appropriate outputs, such as in producing summarisation, reports, and appropriate source codes.

LLMs also act as the bridge that connects people to fully tap into the capabilities of other generative models, such as diffusion models for image, audio, and video generation. Understanding the intricate interplay between these models allows for a deeper appreciation of the unique contributions of LLMs in unlocking the full potential of GenA.I. Most solutions operate through a mechanism that can be broken down into three core steps: (1) user prompt, (2) input processing, and (3) output generation (see Figure 10).

Figure 10: GenA.I. Mechanism



Source: Subject matter expert interviews, Quinlan & Associates analysis

- 1 User Prompt:** A user provides a prompt in natural language—such as a query, a request for specific output, or an instruction, potentially accompanied by media upload or sharing—to GenA.I. application interface powered by an LLM. This step initiates the interaction between the user and the GenA.I. technology.
- 2 Input Processing:** The LLM processes the prompt by converting the input into a semantically rich and coherent numerical representation, expressed as a discrete or continuous representation. Depending on the task, the LLM further interprets the context and intent of the user, preparing the information for the next step in generating the desired output.
- 3 Output Generation:** Based on the processed input, the LLM and/or diffusion models generate the desired output. The LLM ensures that the output is contextually aligned with the user's input, producing relevant results either independently (i.e., text and code) or in

collaboration with diffusion models (i.e., image, audio, and video). The output is then presented in a format that is both easily understandable and accessible to the user.

Through this three-step process, GenA.I. is capable of producing high-quality and easily digestible, user-friendly output without necessitating extensive technical expertise, which allows institutions to optimise their resources, including time and money.

2.3. GenA.I. Development Status

The rapid advancement of GenA.I. has been significantly driven by the availability of vast amounts of training data, increasingly sophisticated model architectures, and the exponential growth in computational power. These models have exhibited remarkable capabilities across a broad spectrum of domains (see Figure 11).

Figure 11: GenA.I. Capability Development

CAPABILITY	TASK	DESCRIPTION	BENCHMARK	HUMAN BASELINE	PERFORMANCE TIMELINE				
					2019	2020	2021	2022	2023
Language & Reasoning	Understanding	Understand the English language and subject knowledge	Massive Multitask Language Understanding ¹	89.8%	32.4%	53.9%	60.0%	75.2%	90.04%
	Factuality & Truthfulness	Ensure factual accuracy and avoid content hallucination	TruthfulQA ² Performance	–			29.5%	–	59.0%
	Mathematical Reasoning	Reason mathematically	MATH ³ Performance	90%			6.9%	51.8% (est.)	84.3%
	Visual Reasoning	Reason across both visual and textual data	Visual Commonsense Reasoning ⁴ Performance	85	66.8	70.8	72	75.6	81.6
	Moral Reasoning	Navigate and reason about ethical principles & moral considerations	Alignment with human judgments ⁵	–	26.6 (GPT-2-XL)	26.6 (GPT-3-curiev1)	–	32.3 (GPT-3.5-davinci)	41.9 (GPT-4)
Coding	Generation	Generate usable code to solve computer science problems	HumanEval ⁶ Performance	–			32.2%	85.1%	96.3%
Image	Generation	Generate images that are indistinguishable from real ones <i>Prompt: a painting of a fox sitting in a field at sunrise in the style of Claude Monet</i>	–	–					
Audio	Speech Recognition	Identify spoken words and convert them into text	Accuracy Rate ⁷ VoxCeleb dataset	–	97.1%	99.4%	99.6%	99.9%	–
Video	Generation	Generate videos from text / images	Difference from reality FVD ⁸ on UCF-101 dataset	–			700	76	36

¹Mass Multitask Language Understanding is a collection of multiple-choice questions spanning academic subjects (e.g., mathematics, philosophy, law, medicine, etc.) to evaluate the capabilities of language models; ²TruthfulQA evaluates whether models generate false answers learnt from imitating human texts; ³MATH is a dataset of challenging competition-level mathematics problems; ⁴Visual Commonsense Reasoning tests how well A.I. systems can reason based on visual data to answer textual questions; ⁵Measured by a discrete agreement metric based on the model's evaluation of a presented situation; ⁶HumanEval is a benchmark for evaluating A.I. systems' coding ability; ⁷Accuracy is measured by 1 - Equal Error Rate; ⁸Fréchet Video Distance is a metric to evaluate the similarity between two video sequences. Lower FVD suggests that the output video is more relevant and thus preferred.

Source: Stanford University - Artificial Intelligence Index Report (2023, 2024), OpenAI, Quinlan & Associates analysis

In natural language processing, GenA.I. has surpassed human baselines in English comprehension and is approaching parity with human performance in areas such as mathematical and verbal reasoning. With an increasing ability for moral reasoning, GenA.I. can engage in coherent conversations, perform complex summarisation tasks, and generate original content, demonstrating its growing versatility.

Speech recognition, a critical component of speech-to-speech GenA.I. use cases, achieved an impressive accuracy of 99.9% by 2022⁹. Although comprehensive evaluation frameworks for generated audio content remain underdeveloped, GenA.I. has already demonstrated the capability of producing lifelike speech and original music, signalling its potential for transformative applications in audio generation.

While GenA.I. has made significant strides in traditional input and output modalities, its scope continues to expand, pushing the boundaries of what was once considered impossible. For instance, in code generation, a leading GenA.I. model successfully solved 96.3% of HumanEval's Python-focused programming problems on the first attempt¹⁰, showcasing its advanced understanding of programming languages and its potential to revolutionise software development.

In the visual domain, various GenA.I. models have emerged that can generate photorealistic images from text prompts, thereby transforming creative workflows with far-reaching implications for product design, visualisation, and the creative arts.

Moreover, by integrating image and audio generation, GenA.I. has developed the capability to animate static images, manipulate videos, and create entirely new clips. Performance metrics such as the Fréchet Video Distance score show an almost 95% reduction, reflecting a significant improvement in the similarity between generated videos and benchmark dataset samples, and indicating the increasing realism of these generated videos.

The continuous advancements in GenA.I. capabilities are poised to reshape the way we interact with and create digital content, with the potential to transform business operations and customer interactions across a multitude of industries, including financial services.

2.4. Conclusion

GenA.I. represents a significant evolution in the capabilities of A.I., extending its reach beyond traditional data processing to include the generation of original outputs. While GenA.I. excels at generating relevant and seemingly novel outputs based on context, the nature of these outputs may be controversial. By design, GenA.I. models predict the likelihood of the next occurrence based on the given context, which includes all previous prompts and generated outputs. As a result, they tend to produce outputs similar to what they were trained on based on the previous data, with intentional randomness introduced to generate more diverse outcomes.

Driven by the advancements in LLMs that unlock the potential of other generative models, GenA.I. is redefining how industries engage with and leverage A.I. technologies. The continuous progress in increasingly sophisticated model architectures, training methodologies, and computational power, is laying the groundwork for widespread application across various sectors.

In the context of financial services, GenA.I. holds the promise of delivering highly tailored customer experiences through personalised financial advice, automated report generation, and dynamic customer engagement strategies. Moreover, its ability to process and analyse vast datasets with unprecedented speed and accuracy presents significant potential for optimising risk management, enhancing fraud detection, and ensuring compliance with regulatory requirements.

⁹ Stanford Institute for Human-Centered Artificial Intelligence. 2023. *2023 AI Index Report*. (https://aiindex.stanford.edu/wp-content/uploads/2023/04/HAI_AI-Index-Report_2023.pdf).

¹⁰ Stanford Institute for Human-Centered Artificial Intelligence. 2024. *2024 AI Index Report*. (https://aiindex.stanford.edu/wp-content/uploads/2024/05/HAI_AI-Index-Report-2024.pdf).

3. GenA.I. Adoption and Impact in the Financial Services Industry

With rapid advancements in its capabilities, GenA.I. is touted to have a profound impact on the financial services industry. With rising demand for enhanced employee productivity and greater customer personalisation, financial institutions in Hong Kong have been leveraging GenA.I. solutions to drive growth and deliver better experiences to their stakeholders.


To assess the current state of GenA.I. adoption among financial institutions, the HKMA conducted an industry-wide survey participated by 137 industry practitioners and in-depth interviews with 16 organisations across the technology, banking, securities, and insurance sectors to gain insights into current and planned GenA.I. use cases.

From these interviews, GenA.I. adoption was broadly categorised into two categories: (1) internal-facing adoption cases for employee empowerment and

(2) external-facing adoption cases for customer empowerment (see Figure 12).

- 1 Internal-facing Adoption Cases:** GenA.I. solutions used to augment employee capabilities by automating a large portion of repetitive and/or low value-adding tasks, allowing staff to produce greater volumes of higher-quality outputs within the same timeframe. These solutions are designed to enhance internal efficiency and can indirectly improve customer engagement through delivering more personalised services.
- 2 External-facing Adoption Cases:** GenA.I. solutions offered by financial institutions that directly service external customers, enabling customers to receive tailored services around the clock without the need for human intervention (e.g., relationship manager (RM), product team, customer support team, etc.).

Figure 12: Types of GenA.I. Adoption

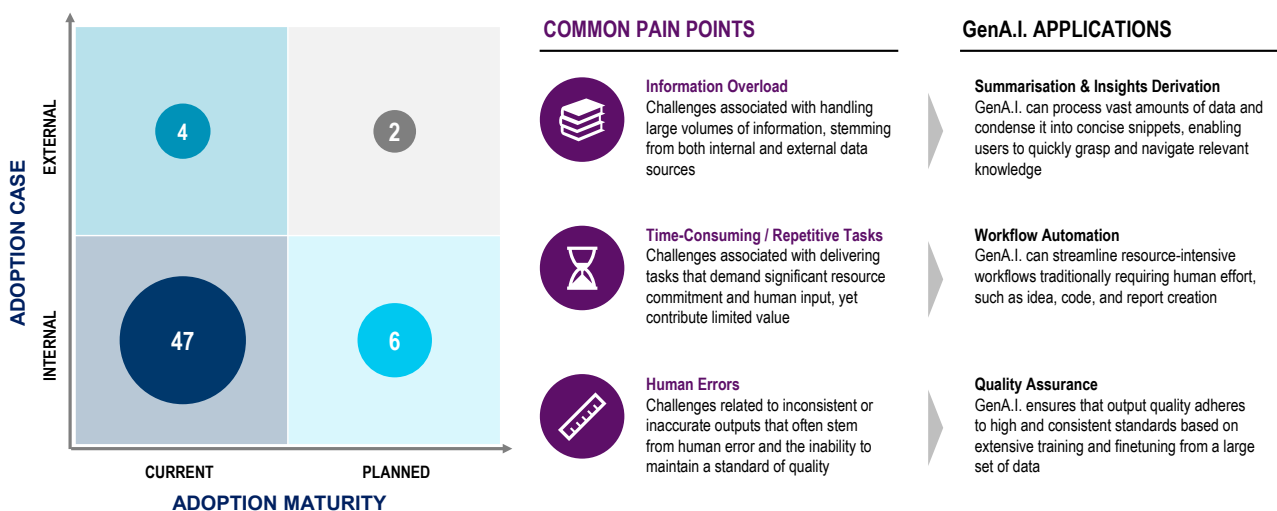


Source: Interview findings, Quinlan & Associates analysis

The institutions that participated in the interview shared a total of 59 GenA.I. use cases, including both current and planned adoption scenarios. Our findings indicate that most organisations are currently focusing on internal-facing use cases, with 47 of the reported cases aimed at supporting employee empowerment (see Figure 13). This emphasis is largely driven by the perception that the risks associated with internal applications are more manageable and controllable.

While some institutions are exploring external-facing applications, adoption has been limited due to uncertainties regarding both technological and regulatory maturity. Most financial institutions we interviewed expressed concerns about whether GenA.I. is sufficiently advanced to consistently deliver reliable and accurate results directly to customers. One technology provider observed that their financial institution clients expect improved technological maturity before adopting GenA.I. solutions and are placing greater emphasis on the explainability of outputs.

Figure 13: GenA.I. Adoption Status



Source: Interview findings, Quinlan & Associates analysis

We observed that financial institutions are adopting GenA.I. to address three key pain points:

1. Information Overload

Financial institutions handle large volumes of information from both internal (e.g., product guidelines, operational manuals, etc.) and external data sources (e.g., third-party research reports, financial disclosures, macroeconomic data, etc.). This often overwhelms employees, particularly new joiners, making it challenging to process and extract valuable insights promptly for decision-making. GenA.I. supports in processing large volumes of data and condensing it into concise, digestible snippets, allowing users to quickly comprehend and navigate relevant information.

2. Time-Consuming / Repetitive Tasks

In addition, financial institutions often commit significant resources and human input to tasks such as idea generation, code generation, outreach message customisation, and report drafting. GenA.I. addresses this by streamlining resource-intensive workflows and automating repetitive and lower-value tasks that require minimal human oversight. By either augmenting or fully replacing these workflows, GenA.I. allows staff to focus on more strategic, high-value activities (see Figure 14).

Figure 14: GenA.I. Workflow Applicability

	Description	Human-in-the-Loop
IDENTIFY NEEDS	Evaluate key objectives, such as organisational goals, regulatory requirements, and customer demands, to pinpoint challenges or opportunities that necessitate actions.	MODERATE Augmentation
GENERATE IDEAS	Brainstorm and develop potential strategies and approaches to effectively address the identified challenges or opportunities.	MODERATE Augmentation
COLLECT DATA	Gather pertinent data (e.g., market reports, customer information, etc.) from both internal and external sources, to support informed decision-making.	LOW Automation
ANALYSE DATA	Apply analytical tools and techniques to process and interpret collected data, identifying significant patterns, trends, and anomalies.	LOW Automation
DERIVE INSIGHTS	Convert data analysis into actionable insights with direct implications for decision-making.	LOW Automation
GENERATE OUTPUT	Create outputs based on findings that align with and effectively address the defined challenges or opportunities.	LOW Automation
REVIEW OUTPUT	Rigorously review outputs for content accuracy, tone consistency with organisational standards, and alignment to industry best practices.	HIGH GenA.I. Not Applicable
TAKE ACTION	Implement decisions (e.g., trade execution, client request handling, etc.) based on the reviewed outputs.	HIGH GenA.I. Not Applicable

Source: Quinlan & Associates analysis

3. Human Errors

Inconsistencies or inaccuracies in outputs resulting from human error present a significant challenge to many organisations. GenA.I. tackles this issue by delivering consistent, high-quality outputs through continuous fine-tuning with vast datasets. This reduces the likelihood of errors and helps ensure a standard of quality across all GenA.I.-generated and revised outputs.

By leveraging GenA.I. for both augmentation and automation, financial institutions can significantly enhance employee productivity while offering greater personalisation for customers, which is critical in maintaining an edge in today's highly competitive financial services landscape.

“

The insurance industry embraces innovations that enhance business operations and customer convenience, bridge protection gaps and deepen financial inclusion. The IA is committed to fostering an environment that is conducive to the development and responsible use of A.I. technologies, including GenA.I., within the insurance market.


Mr Clement Lau, Executive Director, Policy and Legislation, IA

”

Of the 59 cases we examined in our interviews, 8 were general use cases that are both industry- and department-agnostic, such as solutions for taking meeting minutes and code generation. While these agnostic use cases significantly enhance overall workforce productivity, they may not be the most

relevant for financial institutions seeking to improve their specific financial operations. As such, this section examines the 51 use cases specifically relevant to the financial services industry, focusing on their adoption in service fulfilment and operations & risk management (see Figure 15).

Figure 15: GenA.I. Adoption Areas for the Financial Services Space



Source: Interview findings, Quinlan & Associates analysis

- **Service Fulfilment:** Financial institutions seek to help customers fulfil their objectives across a service fulfilment value chain, i.e., (1) customer acquisition, which involves outreach optimisation and content creation, (2) customer monetisation, which involves research and analysis and purchase facilitation, and (3) customer maintenance, which covers customer servicing and ongoing disclosures. Financial institutions are using GenA.I. solutions to deliver a seamless, cohesive, and tailored experience at every stage. Based on the interview findings, they are primarily adopting GenA.I. for customer monetisation and maintenance, with the insurance sector more actively exploring customer acquisition

use cases. This trend is further validated by interview findings from technology providers, revealing that the common GenA.I. adoption cases among the financial institutions they support are focused on customer monetisation and maintenance.

- **Operations & Risk Management:** In this domain, GenA.I. is deployed to optimise back-end internal processes. While it does not directly impact customer-facing functions, such solutions are crucial for driving efficiency and maintaining regulatory compliance. Additionally, GenA.I. solutions are increasingly being used to validate A.I. models themselves and empower institutions to assess vulnerabilities, thereby bolstering overall security and resilience.

3.1. Current Adoption Cases for Service Fulfilment

3.1.1. Customer Acquisition

A customer journey begins with customer acquisition, where financial institutions identify and shortlist target customers, and reach out to—and convert—leads. As part of this process, marketing and sales

teams have to generate relevant marketing material, identify prospective customers, and shortlist targets. However, financial institutions remain cautious about data management and cybersecurity concerns, with most refraining from feeding in externally sourced and sensitive data to their GenA.I. models. As such, the GenA.I. use cases shared by 2 of the 16 institutions we interviewed focused on applications in search engine optimisation (SEO), and 2 focused on generic product marketing and brand building (see Figure 16).

Figure 16: Customer Acquisition Adoption Cases

	COMPANY COUNT ¹	PAIN POINTS			SOLUTION DESCRIPTION	ADOPTION BENEFITS
		Information Overload	Time-Consuming, Repetitive Tasks	Human Errors		
Outreach Optimisation						
SEO	2		✓		• Optimisation of SEO capabilities	• Greater online marketing outreach and efficacy
Content Creation						
Collateral Generation	2		✓		• Creation of marketing content	• Faster content generation for marketing workflows

¹Figures may not add up to the total count presented in Figure 15 due to double counting in the sub-categories

Source: Interview findings, Quinlan & Associates analysis

Search Engine Optimisation

GenA.I. can be employed for various SEO tasks, including identifying trending topics and high-ranking keywords and phrases, by analysing search engine data and online content. Once these insights are gathered, they can be incorporated by GenA.I. to generate multiple content variations, allowing financial institutions to test and subsequently determine the most effective options. For example, a global bank and a global insurer we spoke to reported similar solutions that amplified their SEO endeavours by tailoring their digital marketing content to align with current search trends and interests of target customer segments. This has helped them to streamline processes that would otherwise demand significant time and effort, as well as maximise content reach to the right target audience.

Collateral Generation

From our interviews, we also see financial institutions adopting GenA.I. to create basic marketing and brand promotion materials. These solutions produce outreach content that the marketing team can adapt, saving time on creation from scratch. Two of the three insurance companies we interviewed have empowered their marketing teams with GenA.I. solutions to create customer-facing marketing content. This has helped them to streamline existing workflows and provided their marketing teams with a foothold in the ideation process.

By leveraging GenA.I. for SEO and general marketing content generation, financial institutions can attract new customers more effectively by optimising campaigns in real-time, tailored to market trends, customer preferences, and search engine algorithms.

3.1.2. Customer Monetisation

Once financial institutions acquire a customer base, their focus shifts to enhancing their services—such as offering more personalised solutions—to differentiate themselves, capture cross-selling opportunities, and drive revenue growth. Achieving these goals requires a deep understanding of customer behaviours,

preferences, and needs, which can only be accomplished through comprehensive research and analysis. However, with vast amounts of consumer and market data at their disposal, yet lacking optimal methods to process such information into actionable insights that can support customer outreach efforts, institutions are increasingly turning to GenA.I. solutions (see Figure 17).

Figure 17: Customer Monetisation Adoption Cases

	COMPANY COUNT ¹	PAIN POINTS			SOLUTION DESCRIPTION	ADOPTION BENEFITS
		Information Overload	Time-Consuming, Repetitive Tasks	Human Errors		
Research & Analysis						
External Research	3	✓	✓		• Provision of relevant and current info • Generation of customised insights	• Less time employees spent on research and report generation
Internal Research	8	✓	✓		• Creation of simplified product materials • Provision of guideline summaries	• Better frontline staff preparation for client meetings • Greater employee comprehension of internal documents
Purchase Facilitation						
Needs Identification	1	✓	✓		• Extraction of relevant client data • Identification of client preferences	• More efficient work from frontline staff • Greater customisation for product offerings
Pitch Customisation	2		✓	✓	• Personalisation of outreach material • Generation of marketing collateral	• More tailored messages to engage clients with • More staff focus on creativity and outreach strategy

¹Figures may not add up to the total count presented in Figure 15 due to double counting in the sub-categories

Source: Interview findings, Quinlan & Associates analysis

External Research

Research and analysis are among the most labour-intensive tasks for financial institutions, which may include the preparation of timely market updates, portfolio commentary, and research reports. Analysts and employees spend substantial time gathering, processing, and interpreting large volumes of quantitative and qualitative data. This effort is magnified by the sheer volume of customers and accounts managed, resulting in repetitive analysis and reporting that strains resources.

To alleviate this strain, 3 out of the 16 financial institutions we interviewed are turning to GenA.I. solutions to synthesise information from various sources, such as financial news outlets and market data feeds, and automate the creation of collateral. According to a GenA.I. technology provider, institutions operating in the capital markets sector, in particular, stand to benefit from GenA.I.'s automation capabilities, which enable large-scale data analysis far beyond the capacity of manual efforts. With this capability, GenA.I. can even continuously monitor data sources and update product reports in real-time, ensuring that information remains both up-to-date and relevant. Another global securities firm used GenA.I. to empower its customers to generate keywords to identify suitable stocks for their portfolios.

Internal Research

Financial institutions manage a vast array of financial products and operational procedures, making it challenging for employees to effectively retrieve, understand, and utilise internal documents. GenA.I. solutions tackle this challenge by generating user-friendly summaries and facilitating more efficient document navigation. Our interviews revealed that this application of GenA.I. is the most popular among financial institutions, with 8 out of 16 interviewees having some form of GenA.I. implemented to support internal research.

For example, a major banking institution in Hong Kong leverages GenA.I. to distil internal product documents and guidelines into concise summaries, complete with references to the original sources. This approach makes it easier for employees to understand key information and better prepare them for sales engagements. Another major bank in Hong Kong with a similar solution is looking to expand its GenA.I. capabilities to consolidate all internal information, transforming it into a knowledge bot where employees can ask questions. This would allow employees to extract the specific information they require, bypassing the need to sift through lengthy documents. A global bank utilises a similar solution to help its RMs navigate internal research documents and investment policies to better serve their customers.

Our interview findings also show that GenA.I. applications in research and analysis are highly versatile, supporting a wide range of functions within financial institutions. One global bank we interviewed is using GenA.I. for multilingual support, which is especially useful for operations across multiple jurisdictions, such as Mainland China and Hong Kong (SAR). The bank leverages GenA.I. to process internal policy documents in simplified Chinese and English to answer queries raised by new employees.

By automating and enhancing research and analysis efforts, financial institutions allow employees to focus on higher-value tasks. A global technology provider highlights how this efficiency boost allows employees to be better prepared for critical tasks like loan approvals, supported by insights generated from GenA.I. solutions.



Needs Identification

Once equipped with insights from research and analysis, materials are handed over to client-facing roles such as RMs. However, RMs often face the challenge of turning these materials into actionable insights tailored to each customer they manage. Given customers' diverse risk profiles, including varying preferences, financial goals, and risk tolerances, a one-size-fits-all approach can lead to missed opportunities for hyper-personalised offerings and/or cross-selling opportunities.

To facilitate a more nuanced understanding of client needs, one of the leading banking institutions in Hong Kong utilises GenA.I. to extract and analyse information such as annual reports and financial information about their corporate clients. This application produces tailored memos that RMs can use to engage clients more effectively, offering solutions that are better aligned with each client's specific needs and preferences. The solution also extends to personal banking, allowing RMs to gain insights into their clients' demographics and preferences, enabling them to provide more tailored propositions that were previously challenging to achieve given the sheer scale of personal banking clientele.

Pitch Customisation

Financial institutions often rely on standardised marketing materials and pitches that demand significant time and effort to develop. Such materials are usually designed to be comprehensive yet general to appeal to a wide audience. However, this comes at the cost of creating content with a degree of specificity that will resonate with different clients. Moreover, as clients evolve and their needs change over time, these materials may need to be regularly revisited and updated, demanding additional time and effort.

GenA.I. solutions can provide tailored suggestions for marketing teams and even create customised campaigns and pitches for distinct client segments, with 2 out of 16 financial institutions we interviewed utilising GenA.I. for this purpose. Among them, a global bank uses GenA.I. to translate its marketing materials into different languages and generate conversation prompts for RMs to engage with clients. Another major bank in Hong Kong utilised GenA.I. to generate marketing outreach material for RMs to provide personalised messages to clients. For instance, if a client holds a large deposit balance, the GenA.I. solution may suggest the RM to propose products such as time deposits based on the client's financial profile. This proactive approach to client engagement can help uncover new revenue streams.

By incorporating GenA.I., financial institutions can overcome the challenges of standardisation and lack of differentiation. This ensures that both client interactions and outreach materials are better tailored to meet each client's unique needs. And this hyper-personalised approach ultimately drives higher success rates, allowing financial institutions to capture previously untapped or under-tapped monetisation opportunities.

3.1.3. Customer Maintenance

Beyond monetisation, financial institutions face two ongoing tasks with their existing customers: (1) maintaining client servicing, and (2) ensuring continuous disclosures. As part of this process, they need to address ongoing client needs, including managing enquiries, handling complaints, and providing client reports. However, financial institutions face a multitude of obstacles in conducting these operations (see Figure 18).

Figure 18: Customer Maintenance Adoption Cases

	COMPANY COUNT ¹	PAIN POINTS			SOLUTION DESCRIPTION	ADOPTION BENEFITS
		Information Overload	Time-Consuming, Repetitive Tasks	Human Errors		
Customer Servicing						
Enquiries Management	7	✓	✓	✓	• Generation of personalised responses • Assistance and monitoring of interactions	• More consistent and faster responses for clients • Greater control over service quality
Complaint Analysis	1		✓		• Analysis of customer complaints • Automatic assessment of staff responses	• Faster complaint analysis times • Greater efficiency in processing
Ongoing Disclosures						
Customer Reporting	1		✓		• Notification of client reporting dates • Generation of draft reports for staff	• Less frontline staff workload on reporting • Faster report generation process

¹Figures may not add up to the total count presented in Figure 15 due to double counting in the sub-categories

Source: Interview findings, Quinlan & Associates analysis

Enquiries Management

Financial institutions need robust enquiry systems, given that the public often lacks an understanding of financial products and requires strong aftersales support. However, these systems face several limitations, namely: (1) addressing the sheer volume of enquiries can hinder response capabilities (e.g., delays), (2) ensuring consistency in service delivery can be challenging, and (3) monitoring overall service quality can be problematic.

To manage these issues, 7 out of 16 financial institutions we interviewed have enhanced their client servicing using GenA.I., the second most popular use case of GenA.I. we identified. By leveraging GenA.I.'s workflow automation features, query systems can generate rapid personalised responses based on customer data. This not only improves the efficiency of customer engagement but also enhances customer retention through more targeted messaging. For instance, one global securities firm reported a 33% reduction in the time needed to draft replies after implementing this solution. Another leading technology provider noted that its insurance client reduced post-call work (e.g., inputting the customer's enquiries) by 30%, allowing staff to focus on determining the best next steps without needing to manually input customer information.

Additionally, several organisations we spoke to pointed to the ability of their GenA.I. solutions to deliver consistency in frontline messaging. One interviewee noted the difficulty of having their staff consistently produce standardised content in a professional tone. To alleviate this issue, it has adopted GenA.I. solutions to create initial draft replies that its employees can simply review and finalise before sending to address customer queries and complaints. This approach has significantly reduced response variance among employees, ensuring that all communications maintain a professional tone and consistent service quality.

Similarly, a global bank we spoke to highlighted that its GenA.I. solution ensures customer complaint responses are consistently delivered in a professional tone and with high-quality service, aligning with its brand image. Another global insurer utilised a similar solution to address customer enquiries, which led to a more than three-fold increase in cases addressed per day.

Some of the interviewees we spoke to have also capitalised on GenA.I. to monitor customer servicing, with a focus on ensuring compliance and assessing service quality. In these solutions, GenA.I. automatically identifies service gaps that may be overlooked by human error and highlights areas for improvement in customer maintenance. These solutions enhance operational efficiency by automatically recognising errors and fostering a culture of continuous quality improvement. One global insurer observed that implementing this GenA.I. solution enhanced its ability to identify issues and systematically pinpoint areas for improvement in its call interactions.

Complaint Analysis

To complement their client servicing capabilities, financial institutions need procedures to handle and escalate complaints. However, this is often easier said than done due to the substantial resources required to understand and investigate complaints. Customer Success staff or related departments must painstakingly—and carefully—listen to and comprehend thousands of complaints while simultaneously conducting their research.

To enhance their complaint analysis and handling capabilities, one financial institution we interviewed has begun to leverage GenA.I. workflow automation features: their solution automates the parsing of complaints, enabling them to investigate and take remedial actions in a more structured manner. Similarly, a global securities firm we interviewed stated that its GenA.I. solution has delivered a 50% reduction in the time taken for complaint call analysis.

Customer Reporting

To effectively manage accounts, whether they are investment, banking, or insurance policy accounts, financial institutions routinely generate and share periodic reports or statements with end-users, detailing historical transactions and activities. While automation has streamlined the generation of these statements to some degree, the process of producing

detailed reviews can still be time-consuming and resource-intensive. Often, these reports require the integration of relevant commentary and nuanced analysis that current automation technologies cannot fully address.

One global bank that we interviewed has adopted GenA.I. solutions to create detailed, contextually relevant reports for customers by analysing transaction data, identifying key trends, and crafting tailored commentary. This includes streamlining the review process by dynamically adjusting content to reflect the most updated data available. More specifically, their GenA.I. solution helps to draft portfolio reviews that require only modifications by RMs, allowing RMs to focus more on strategic, value-adding tasks.

The integration of GenA.I. into financial institutions' customer maintenance processes enhances their ability to meet customer needs in enquiries management, complaint analysis, and customer reporting. These solutions efficiently summarise vast amounts of information, automate repetitive workflows, and accurately identify inaccuracies and errors, enabling institutions to strengthen customer relationships at scale.

3.2. Current Adoption Cases for Operations & Risk Management

Operations are the foundation of financial institutions, with risk management serving as a critical mechanism for safeguarding customer trust and ensuring regulatory compliance. In an increasingly complex and dynamic competitive environment, financial institutions are striving to enhance their operational efficiency and resilience by turning to GenA.I. technologies. These solutions offer the ability to automate labour-intensive processes, especially the analysis of large datasets, which is useful for detecting patterns and anomalies that support proactive risk management (see Figure 19).

Figure 19: Operations & Risk Management Adoption Cases

	COMPANY COUNT ¹	PAIN POINTS			SOLUTION DESCRIPTION	ADOPTION BENEFITS
		Information Overload	Time-Consuming, Repetitive Tasks	Human Errors		
Risk Profiling						
Customer Due Diligence	4	✓	✓		- Automation of due diligence checks - Summarisation of client documentation	- Quicker method to locate relevant information - Greater overall efficiency
Regulatory Monitoring						
Compliance Review	2	✓		✓	- Summarisation of policy documentation - Identification of policy compliance gaps	- Greater employee understanding of policy documents - Greater risk mitigation for external compliance
Technology Risk Management						
A.I. Model Validation	2		✓		- Summarisation of A.I. model assessment - Analysis of A.I. capabilities and suitability	- Faster A.I. model assessment - More efficient operations that are compliant
A.I. Model Compliance Check	2	✓	✓		- Assurance of compliant model behaviour - Generation of compliant responses	- More efficient and compliant operations with GenA.I. - Greater confidence in adhering to standards
Cybersecurity	2	✓	✓		Automation of cybersecurity alert processing	- More efficient prioritisation of alert by staff - Faster cybersecurity reporting process

¹Figures may not add up to the total count presented in Figure 15 due to double counting in the sub-categories

Source: Interview findings, Quinlan & Associates analysis

Customer Due Diligence

Financial institutions must evaluate the financial health and risk profile of new customers to ensure compliance with KYC and AML regulations. However, the customer due diligence process involves analysing vast amounts of data within submitted documents, taking up a significant amount of time and effort.

To streamline this essential operation, 4 of the 16 organisations interviewed have implemented GenA.I. solutions to enhance risk assessments, enabling employees to profile client risks more efficiently and with greater accuracy. This includes rapidly summarising customer information to quickly grasp their situation and locate relevant details for further investigation. By automating the document review process, three interviewees highlighted their ability to arrive at faster and more accurate due diligence decisions, with a global credit rating agency citing a 27% reduction in time needed for credit analysis tasks. One regional bank echoed this sentiment, stating that GenA.I.'s summarisation

capabilities and automated checking processes allowed its employees to focus on more complex tasks. Two global banks also reported improved decision-making efficiency and enhanced customer understanding with the support of GenA.I. solutions.

While other interviewed organisations have yet to adopt such solutions, there is substantial potential for securities firms to leverage GenA.I. solutions to identify unusual trading activities while examining new customer profiles during KYC checks. Insurance firms may also benefit from producing preliminary reports, based on a customer's previous claims history, to detect potential fraudulent claims more rapidly. More broadly, leveraging GenA.I. solutions not only enables financial institutions to expedite the onboarding process for consumers but also facilitates the creation of more accurate customer profiles, allowing for the delivery of personalised products and services.

Compliance Review

As the regulatory landscape evolves, with updates to regulations, circulars, and guidelines, financial institutions must maintain continuous compliance monitoring. This practice can be time-consuming, requiring compliance personnel to regularly sift through lengthy regulatory documents and extensive internal policy manuals. When addressing new regulatory requirements or conducting general reviews, the sheer volume of documentation creates challenges in pinpointing and reconciling information.

To support the compliance review process, two financial institutions are using GenA.I. to enhance staff capabilities in responding to new policies. Their solutions provide clear, concise breakdowns of policy documents, allowing employees to swiftly understand key regulatory changes. For example, a global bank utilises GenA.I. to allow employees to search through documents and quickly locate the most relevant regulations, policies, and procedures that need to be followed.

Another regional bank leverages similar technology to identify gaps between current operational policies and new or updated regulations, facilitating faster responses and reducing manual errors. As they noted, adopting GenA.I. in this capacity has made the bank more efficient at identifying compliance gaps and taking swift action.

While adoption has thus far been observed primarily in banking institutions, securities and insurance firms can similarly benefit from similar applications for regulatory compliance. Overall, incorporating GenA.I. into compliance reviews enables institutions to swiftly adapt to regulatory changes and address potential compliance shortfalls.

A.I. Model Validation

With the growing exploration of A.I. and GenA.I., financial institutions must properly assess, validate, and verify the A.I. models' output to ensure they effectively fulfil their intended purposes. The model validation process is resource-intensive, requiring

the analysis of large datasets and the evaluation of various models. Once validation is complete, institutions must compile comprehensive reports documenting their findings, a process can take several months for a single use case, as cited by a banking institution we interviewed.

To support their model validation endeavours, a regional insurer has adopted GenA.I. to assess different A.I. models, helping users to efficiently choose the model best suited to their specific needs. Additionally, a global bank utilises GenA.I. to automate the drafting of validation reports, summarising key insights and findings. This has not only reduced the workload associated with report preparation but also improved employees' ability to interpret insights more effectively due to the synthesised nature of the output.

By leveraging GenA.I. solutions, financial institutions can significantly reduce the time and effort required for model validation, accelerating the deployment timelines of models in practical application.

A.I. Model Compliance Check

After validating the model functionalities, financial institutions must ensure that these models comply with relevant regulations. Off-the-shelf solutions, however, often present challenges as they may not fully align with financial industry standards, particularly in data privacy and intellectual property compliance.

To address these challenges, two of the institutions interviewed have developed tailored GenA.I. solutions, enabling them to maintain compliance when using code generated by their GenA.I. systems. A global bank utilises its GenA.I. solution to reference its proprietary approved code repository, which contains millions of lines of tested code, allowing developers to generate and review code while ensuring compliance. Similarly, a securities firm has implemented a customised version of ChatGPT, incorporating additional features to enhance interaction with content, allowing its employees to use ChatGPT in a secure and compliant manner.

These tailored GenA.I. solutions allow financial institutions to innovate while maintaining the compliance standards essential to their operations. By doing so, institutions can harness the productivity gains of GenA.I. without compromising their regulatory obligations.

Cybersecurity

In today's digital landscape, financial institutions are often inundated with an overwhelming volume of cybersecurity alerts, potentially numbering in the thousands each day. This influx of notifications presents a significant operational challenge, particularly for institutions with limited resources or insufficient cybersecurity expertise. The sheer volume of alerts increases the likelihood that critical vulnerabilities may be overlooked, thereby heightening the risk of security breaches and prolonging response times.

To mitigate these challenges, two financial institutions interviewed have adopted GenA.I. technologies to optimise the analysis of cybersecurity notifications and reports. By automating the extraction of core insights from the alerts, these financial institutions

have streamlined their cybersecurity operation, significantly reducing the time spent on manual investigations and enabling more efficient prioritisation of issues. For instance, a multinational securities firm utilises GenA.I. copilots to filter alerts based on priority so that employees can focus their attention on the most important cybersecurity issues, accelerating their internal cybersecurity processes.

In addition, the GenA.I. solution can supplement internal analytics by leveraging external datasets to tackle cybersecurity issues. A notable application was observed in an insurance firm, which utilises GenA.I. solutions to generate reports for external attack service management and vulnerability assessments. Users can interact with reports dynamically, asking specific questions about vulnerabilities and quickly accessing essential information.



By harnessing targeted cybersecurity analysis powered by GenA.I., financial institutions can significantly enhance their decision-making processes and expedite their responses to potential threats, cultivating a proactive approach to risk management. This helps to strengthen the overall security posture and foster a culture of vigilance and resilience within financial institutions, better preparing them to navigate an increasingly complex digital landscape.

3.3. Planned Adoption Cases

As GenA.I. is still in the early stages of exploration for many financial institutions, and its advancements often outpace deployment timelines, several use cases are still being examined by the organisations we interviewed.

Two prominent themes of planned adoption cases emerged from the interviews, including (1) external-facing applications, such as customer-facing chatbots and (2) non-text processing applications in more creative domains, like avatar generation. While such use cases are promising, all interview participants recognise that their implementation will depend on greater technological reliability and regulatory maturity. Consequently, they were viewed more as long-term aspirational goals.

3.3.1. External-Facing Applications

The majority of current GenA.I. use cases are internal-facing, primarily focused on employee empowerment, as the presence of a human in the loop helps manage risks. While external-facing applications remain limited, financial institutions are showing significant interest in adopting GenA.I. for these purposes. As a result, many institutions are tentatively exploring these solutions while awaiting clearer regulations.

Customer-facing Chatbots

One notable solution is customer-facing chatbots, which enhance self-service capabilities by providing tailored, context-aware responses. Currently, many internal-facing use cases are a form of query-based interaction that retrieves and summarises data from knowledge bases. Given the success of these solutions, it is only natural that financial institutions

are thinking about extending these functionalities to a broader user base, such as consumers, once regulations mature.

These capabilities streamline client interactions, offering faster and more timely service. For example, customers can directly file claims, inquire about policy coverage, and check claims processing times online at any time, all without needing to contact a call centre. The ability to access high-quality service at any time, unrestricted by traditional working hours, significantly boosts client satisfaction and overall experience.

Embedded Tools

In addition to customer-facing chatbots, financial institutions also see the potential of GenA.I. applications serving as embedded tools to deepen client engagement. These solutions can offer advanced analysis, summarisation, and predictive functionalities directly to clients, facilitating better financial decisions. For instance, they can provide securities clients with valuable investment insights, assist banking clients in making informed product choices, and help insurance customers select the most suitable coverage options.

By discreetly integrating client information within prompts, these GenA.I. solutions can automatically recommend more appropriate products and services. For example, a securities brokerage firm could use GenA.I. to recommend portfolio adjustments based on news from user-preferred sources that were shortlisted by the users and individual risk tolerances. Once the GenA.I. tool detects relevant market updates, the embedded tool could immediately notify the user of corresponding trading opportunities.

While many of these applications are technically feasible, most financial institutions have yet to implement these solutions due to the lack of regulatory maturity. In particular, certain system-generated suggestions may be classified as regulated activities, such as investment advice. Additionally, the outputs of these systems are often challenging to control, making it crucial to establish a clear liability framework. The interviewed institutions recognise that these applications promise to exponentially enhance user experience and improve operational efficiency. We anticipate that these institutions are poised to adopt more external-facing solutions once regulations further mature.

3.3.2. Non-text Processing Applications

While current GenA.I. use cases primarily focus on text processing, the technology's capabilities extend to a wide range of formats, including images, audio, and video. By harnessing these capabilities, the potential applications of GenA.I. can broaden significantly, transforming essential processes such as client support and customer engagement.

Non-text Analysis

Traditional client interactions can be suboptimal, as RMs and customer support teams often miss critical information and cues conveyed by consumers. GenA.I.'s ability to process various input formats allows for the examination of non-verbal cues, including body language, facial expressions, and tone of voice. These subtle aspects of communication that are often overlooked in conventional interactions can serve as powerful tools to enhance the understanding of client sentiments and reactions.

The capability of GenA.I. can manifest in digital avatars, which interpret a customer's emotional state and identify signs of confusion, frustration, or satisfaction in real-time. This technology ensures that no information is overlooked, capturing even the smallest cues. Consequently, support teams are empowered to tailor their responses more effectively based on the context of each interaction. For example, the GenA.I.-powered digital avatar may detect anxiousness in the consumer and prompt an RM to provide additional information in a reassuring tone to address their concerns.

The purpose of such solutions is not to replace human interactions with a fully digital, automated proposition, as many consumers still value a human touch. Rather, GenA.I. can empower the agents who are interacting with consumers. Ultimately, the integration of GenA.I. holds the potential to transform the landscape of consumer engagement, enabling financial institutions to move beyond transactional relationships to cultivate more meaningful connections with consumers.

Non-text Generation

As their internal operations are not limited to text-based interactions, financial institutions are also exploring the adoption of GenA.I. in the realm of non-text generation. By utilising different GenA.I. models, like diffusion models, financial institutions can harness the technology's image, audio, and video generation abilities, which can be utilised to automate workflow and enhance customer engagement.

With GenA.I.'s image and video creation capabilities, financial institutions can enhance their marketing campaigns with customised posters and other media. Automating the design process allows marketing departments to produce high-quality content efficiently, enabling rapid responses to trends and promotional opportunities, saving time and reducing associated costs. Financial institutions also seek to use similar capabilities to visualise complex financial information for client-facing educational materials. Clients can interact with GenA.I. solutions to create infographics from key financial data to better understand their financial situation, supporting more informed decision-making.

Financial institutions also show interest in leveraging GenA.I.'s audio generation capabilities for automated client servicing. One bank aims to use GenA.I. for automated calls to check in with clients, remind them of debt repayments, and understand delays. This would free employees to focus on complex tasks while ensuring clients receive timely notifications, improving the efficiency of debt collection. Another potential application is using GenA.I. solutions to handle simple enquiries through hotlines, generating appropriate responses, to address issues around the clock.

Together, these capabilities significantly enhance customer engagement. By providing relevant and easily digestible information and 24/7 support, financial institutions can foster stronger relationships with clients. This improved communication leads to higher levels of satisfaction and loyalty, ultimately driving business growth.

Overall, financial institutions are considering several exciting applications, especially for external-facing applications and non-text data. The institutions we interviewed have expressed enthusiasm about adopting these solutions, provided they have adequate resources, robust technological knowledge, and regulatory certainty.

3.4. Adoption Hurdles

Financial institutions recognise that adopting A.I. solutions will significantly impact both service fulfilment and operations & risk management, with approximately 80% of industry practitioners sharing this view in the 2023 HKMA Fintech Adoption Study. Consequently, we have observed considerable efforts

by financial institutions to advance A.I. adoption, with GenA.I. becoming a prominent focus alongside technologies like robotic process automation and optical character recognition. However, the survey also identified three major adoption hurdles that could be addressed to promote broader GenA.I. adoption, which were echoed by the financial institutions we interviewed (see Figure 20).

Figure 20: Key Adoption Hurdles


Expected Level of Impact of A.I. Adoption

% Moderate / High Impact, 2023, n = 137



Key Adoption Hurdles

% Moderate / Major Challenge, 2023, n = 137



Source: HKMA Fintech Adoption Study (2023), Quinlan & Associates analysis

1 Compliance / Risk Concerns: Compliance and risk management are critical factors in the decision-making process for GenA.I. adoption. Financial institutions are particularly cautious about potential risks associated with GenA.I. solutions, especially the risk of information leakage. For example, one of the insurers emphasised the need for caution when inputting sensitive data into A.I. systems to avoid potential breaches. Solution providers are also acutely aware of the ethical implications and are adopting responsible data management practices to mitigate these risks. A global bank mentioned that a well-defined framework with a technology-neutral, risk-based, and outcome-focused approach that can strike a balance

between addressing risks and encouraging innovation would be helpful. A technology provider also suggested that a regulatory classification of use cases based on their level of risk would streamline the adoption of GenA.I., as even basic solutions may currently face the same level of scrutiny as more advanced ones.

2 Lack of Practical Knowledge: Due to the recent emergence of GenA.I., many traditional financial institutions may lack the knowledge and expertise to effectively train and implement these solutions using appropriate data. This knowledge gap underscores the need for suggested practice guidance to ensure that GenA.I. solutions are both effective and useful,

which is essential for broader adoption. A banking institution expressed that it would be helpful for regulators to provide training on the adoption of GenA.I. in the banking sector. Similarly, a global securities firm noted the importance of educational programmes to nurture a skilled workforce. Aside from educational or training programmes, use case sharing from other banks was noted by a banking institution as a valuable way to gain insights into navigating the regulatory requirements for GenA.I. adoption.

3 Extensive Resource Requirements:

Implementing GenA.I. solutions requires considerable resource commitment. Financial institutions must be prepared for substantial upfront costs to ensure that the solutions are customised to their specific needs and can deliver a sufficient return on investment over time. Additionally, the development process involves extensive research, data collection, processing, model training, testing, validation, and fine-tuning. As a result, financial institutions need to allocate ample time and resources to GenA.I. adoption projects before they can be effectively deployed in real-world scenarios.

Given these challenges, financial institutions are adopting a cautious approach to GenA.I., which may result in slower implementation timelines than initially anticipated. To facilitate the adoption process, the subsequent sections of this report aim to address these concerns: (1) Section 4 on regulatory principles provides further clarity on compliance and risk management, (2) Section 5 on adoption considerations offers practical guidance across the pre-deployment, deployment, and post-deployment phases of GenA.I., and (3) Section 6 on facilitation initiatives explores how various support schemes can alleviate business-specific deployment concerns.

“

The MPFA recognises the potential of A.I. or GenA.I. to enhance the management of retirement benefits and improve service delivery in the MPF industry. The MPFA will continue to collaborate with other financial regulators and the MPF industry to foster the healthy development of A.I. technologies within the financial services sector.

Mr Kenneth Chan, Executive Director (Members and Supervision), MPFA

”

4. Regulatory Principles



The impact of regulations on the adoption of new technologies cannot be overstated. Thoughtfully crafted regulations can foster innovation and ensure that technologies flourish by providing clear guidelines, protecting users, and establishing trust. Conversely, restrictive or outdated regulations can stifle innovation, preventing new technologies from reaching their full potential.

While A.I. has long been capable of generating outputs, GenA.I. takes this capability to new heights by making it more accessible and user-friendly. This evolution allows a broader range of end-users to leverage A.I.'s generative capabilities, either to replace or augment their daily tasks. The breadth and depth of data used to power these models, combined with the sophistication of the applications built on top of them, are unprecedented.

A.I. has consistently been at the forefront of the global technological agenda, with jurisdictions worldwide recognising its transformative potential across various sectors, especially with the rise of

GenA.I. In response to these advancements, regulators have acknowledged the need to update and refine their approaches to A.I.-related regulations that are broadly also applicable to GenA.I., while addressing GenA.I.-specific novel risks, such as hallucination (i.e., generating inaccurate or fabricated information) and unfair biases (i.e., producing outputs that reflect systemic prejudices or discrimination).

Over the past decade, both before and after the introduction of GenA.I., there has been a concerted effort to create regulatory environments that promote the responsible adoption of the technology while safeguarding end-user interests. Many jurisdictions have provided supervisory guidance based on a set of principles, typically organised into two categories, to ensure the ethical and responsible use of A.I.

- 1 **Common Principles:** Widely recognised principles that are continuously refined by most jurisdictions.
- 2 **Additional Principles:** More specific principles that are emphasised by certain jurisdictions based on their unique priorities (see Figure 21).

Figure 21: Key Regulatory Principles of A.I.



Source: Regulatory disclosures, Quinlan & Associates analysis

Common Principles

These foundational principles are frequently discussed, updated, and incorporated into the A.I. regulatory and/or supervisory frameworks of many jurisdictions. They form the bedrock of responsible A.I. usage and are essential for maintaining public trust and safeguarding end-users' rights. The common principles include:

- 1 Governance & Accountability:** Establishing a clear framework for governance and accountability is essential for the responsible deployment of A.I. systems. This includes assigning clear responsibilities to stakeholders within the organisational structure and ensuring an appropriate level of explainability to help them understand the A.I. system's mechanisms and outputs. These measures mitigate unintended consequences, ensure effective oversight, and address any issues that arise promptly and efficiently.
- 2 Fairness:** Ensuring fairness in A.I. systems is a critical concern. These systems should be designed and implemented to ensure equal and inclusive treatment of all users, preventing bias and avoiding discrimination against any particular group. This principle requires active measures to detect and mitigate potential biases. Fairness also demands that A.I. systems be rigorously tested and evaluated to safeguard equity in their outputs, ensuring that no individual or group is unfairly disadvantaged by the technology.
- 3 Data Privacy & Protection:** Maintaining robust data privacy and protection measures are required for safeguarding user information in A.I. systems. This requires implementing strong security measures, including cybersecurity protocols, to protect personal data and defend against unauthorised access and data breaches. It is equally important to seek explicit user consent for data usage and clearly communicate the intended use of that data. Organisations must comply with existing data privacy laws and implement a robust data protection protocol that ensures sensitive information is securely stored, enabling users to retain control over how their personal data is used in A.I. processes.

- 4 Transparency & Disclosure:** Upholding the principle of transparency and disclosure is crucial for building trust and ensuring that users understand how A.I. systems operate. This requires providing users with clear and sufficient information about the use of A.I. in the services they receive, protecting their rights and ensuring informed consent. Disclosures should include details about the presence of A.I. across the user journey, its functionality, and the decision-making processes involved.

Additional Principles

These principles, while also important, are emphasised in jurisdictions with specific priorities and reflect their unique regulatory environments and strategic objectives. The additional principles include:

- 5 Reliability:** Ensuring the reliability of A.I. systems is needed for their operation. This involves building A.I. systems that are resilient and capable of delivering accurate and dependable results consistently. Operating the system with resilience and using high-quality, diverse training data is crucial to reducing the risk of errors and biases, thereby enhancing user confidence in the technology. Reliability ensures that A.I. systems perform as expected under various conditions and that their outputs remain trustworthy over time. Jurisdictions prioritising reliability often implement rigorous testing and validation procedures to meet high standards of performance and dependability.
- 6 Sustainability:** Actively monitoring and managing the environmental and social impacts of A.I. systems is becoming increasingly important as the technology evolves. This involves efforts to minimise the environmental footprint of A.I. systems, which often demand significant computational resources for training and operation. Promoting the responsible use of A.I. to address global challenges such as climate change, resource management, and social equity is also crucial. Jurisdictions that prioritise sustainability are more likely to implement policies that encourage the development of energy-efficient technologies and the use of A.I. in addressing environmental and societal issues.

Having outlined the key regulatory principles, we will now examine in greater detail how the four common principles are being addressed across six jurisdictions, including (1) the EU, (2) Hong Kong (SAR), (3) Mainland China, (4) Singapore, (5) the U.K., and (6) the US. Each jurisdiction has taken unique approaches to embedding these principles into their regulatory frameworks, reflecting their respective priorities, legal environments, and strategic goals for A.I. innovation and adoption.

For the Governance & Accountability, Fairness, and Data Privacy & Protection principles, we will explore the regulatory updates at the jurisdiction level, as these principles are well-defined and applicable across industries adopting the technology. For the Transparency & Disclosure principle, we will investigate the updates at the industry level, as this principle is often tailored to meet the specific needs of financial market operations.

In the case of the US, however, much of the regulatory developments surrounding A.I. take place at the state level. As such, overarching guidance across these principles is provided by the National Institute of Standards and Technology (NIST)¹¹. NIST's recent framework, titled '*Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*'¹², developed in alignment with President Biden's 2023 '*Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence*'¹³, offers a comprehensive set of

standards and guidelines. Although compliance with the framework is voluntary, it plays a crucial role in helping organisations, individuals, and society at large to manage the risks associated with A.I. deployment more effectively.

4.1. Governance & Accountability

GenA.I. is a transformative tool that can streamline a broad range of tasks, from back-end employee support to front-end customer engagement. However, despite its powerful capabilities, it is not immune to errors or unintended consequences. The responsibility for any issues that arise from the use of GenA.I. ultimately rests not with the technology itself but with the humans (i.e., internal staff) who develop, implement, and oversee these systems. An important component of accountability is to understand the rationale behind GenA.I.-assisted decisions and how the solutions arrive at their outputs.

To address these challenges, regulators across the globe are issuing guidelines that define the roles and responsibilities related to GenA.I. solution providers and setting expectations for interpreting GenA.I. solutions' mechanisms and outputs. Such measures are essential for maintaining accountability and explainability, ensuring public trust, and safeguarding the rights of end-users (see Figure 22).

¹¹ The NIST is a non-regulatory government agency of the US Department of Commerce to promote US innovation and industrial competitiveness.

¹² NIST. 2024. *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. (<https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>).

¹³ The White House. 2023. *Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence*. (<https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>).



Figure 22: Governance & Accountability

OBJECTIVE

Allocate the responsibility clearly to stakeholders and maintain a level of explainability for the A.I. solutions

JURISDICTION	RELEVANT EXCERPTS
European Union European Parliament and the Council of the European Union • <i>Artificial Intelligence Act (2024¹)</i>	- Providers of high-risk AI systems shall put a quality management system in place that ... shall include at least ... an accountability framework setting out the responsibilities of the management and other staff with regard to all the aspects listed in this paragraph. - High-risk AI systems shall be designed and developed in such a way as to ensure that their operation is sufficiently transparent to enable deployers to interpret a system's output and use it appropriately.
Hong Kong (SAR) Digital Policy Office of the Hong Kong Government • <i>Ethical Artificial Intelligence Framework (2024)</i>	- The three lines of defence is a well-established governance concept in many organisations. Organisations are responsible for the moral implications of their use and misuse of AI applications. There should also be a clearly identifiable accountable party, be it an individual or an organisational entity (e.g., the AI solution provider). - Two approaches can be taken to generate human-readable explanations... built-in interpretation... (and) post-hoc interpretation.
Mainland China National New Generation Artificial Intelligence Governance Specialist Committee • <i>Ethical Norms for New Generation Artificial Intelligence (2021)</i>	... improve... explainability, comprehensibility... to gradually achieve verifiability and auditability. ... humans are the ultimate responsible parties. Clarify the responsibilities of stakeholders ... establish an artificial intelligence accountability mechanism... Clarify the responsibilities and power boundaries of Artificial Intelligence-related management activities.
Singapore Infocomm Media Development Authority and Personal Data Protection Commission • <i>Model Artificial Intelligence Governance Framework Second Edition (2020)</i>	- Organisations using AI in decision-making should ensure that the decision-making process is explainable... (and) that their use or application of AI is undertaken in a manner that reflects the objectives of these principles as far as possible. - Personnel and/or departments having internal AI governance functions should be fully aware of their roles and responsibilities, be properly trained, and be provided with the resources and guidance needed for them to discharge their duties.
United Kingdom Information Commissioner's Office and The Alan Turing Institute • <i>Explaining Decisions made with AI (2022)</i> Competition & Markets Authority • <i>CMA (Competition & Markets Authority) AI Strategic Update (2024)</i>	- Actively consider and make justified choices about how to design and deploy AI models that are appropriately explainable to individuals. Take steps to prove ... that they are present in the design and deployment of the models themselves. Show that you provided explanations to individuals. - Foundational model developers and deployers are accountable for FM outputs... (and) take responsibility for what they control in the value chain and take positive action necessary to ensure consumers are adequately protected.
United States National Institute of Standards and Technology • <i>Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (2024)</i>	- Accordingly, GAI (GenA.I.) may call for different levels of oversight from AI Actors or different human-AI configurations in order to manage their risks effectively. Governance tools and protocols that are applied to other types of AI systems can be applied to GAI systems. - Implement interpretability and explainability methods to evaluate GAI system decisions and verify alignment with intended purpose.

¹Rules entered into force in August 2024, and provisions will apply 12-24 months after
 Source: Regulatory disclosures, Quinlan & Associates analysis



- EU:** Recognising the rapid development of A.I. solutions, the '*Artificial Intelligence Act*'¹⁴ was established as the world's first comprehensive regulatory framework for this technology. To ensure appropriate interpretation and deployment, Article 17 of the Act requires an accountability framework that defines the responsibilities of management and staff concerning various aspects of high-risk A.I. systems, including technical specifications, data management, and post-market monitoring. Article 13 also mandates that instructions for use provide clear information about the system's technical capabilities and characteristics, as well as guidance for interpreting and appropriately using the A.I. system's outputs.
- Hong Kong (SAR):** To guide the applications of A.I., the Digital Policy Office of the Hong Kong Government introduced '*Ethical Artificial Intelligence Framework*'¹⁵. The framework's accountability principle states that organisations are responsible for the moral implications of using and misusing A.I. applications and proposed three lines of defence in A.I. governance that institutions can look to set up. It also suggested two approaches to achieve explainability: built-in interpretation and post-hoc interpretation. Built-in interpretation provides transparency by showing the exact steps to arrive at a decision, making it clear to end-users, while post-hoc interpretation may require the use of visual tools or examples to explain a model's behaviour after decisions are made. The Framework encourages other organisations to adopt its principles as a reference for their own A.I. initiatives.
- Mainland China:** In 2021, China introduced '*Ethical Norms for New Generation Artificial Intelligence*'¹⁶, marking a significant step towards ethical A.I. solutions. It expects A.I. systems to be more comprehensible and explainable, gradually achieving verifiability and auditability. Moreover, it underscores that

¹⁴ European Parliament and the Council of the European Union. 2024. *Artificial Intelligence Act*. (<https://ai-act-law.eu/>).

¹⁵ Digital Policy Office of the Hong Kong Government. 2024. *Ethical Artificial Intelligence Framework*. (https://www.digitalpolicy.gov.hk/en/our_work/data_governance/policies_standards/ethical_ai_framework/).

¹⁶ National New Generation Artificial Intelligence Governance Specialist Committee. 2021. *Ethical Norms for New Generation Artificial Intelligence*. (https://www.most.gov.cn/kjbgz/202109/t20210926_177063.html/).

humans are the ultimate responsible parties. Stakeholders' responsibilities and powers should be clearly delineated to standardise the conditions and procedures for the execution of rights to use, manage, and conduct other activities. In particular, the responsibilities of A.I.-related management activities should be clarified.

- **Singapore:** Addressing public concerns about A.I. solutions, Singapore's Infocomm Media Development Authority (IMDA) and Personal Data Protection Commission developed the '*Model Artificial Intelligence Governance Framework Second Edition*'¹⁷ to take a structured A.I. governance approach, incorporating feedback from various stakeholders and experiences from organisations in the second edition. One overarching principle's key component is explainability, requiring organisations to ensure the process is explainable and align their A.I. applications with established principles if they use A.I. in the decision-making process. Personnel or departments responsible for internal A.I. governance should clearly understand their roles and responsibilities, and they must be equipped with the necessary resources and guidance to fulfil their duties effectively.
- **U.K.:** As organisations employ A.I. to support or make decisions about individuals, the Information Commissioner's Office and The Alan Turing Institute issued a co-badged guidance '*Explaining Decisions made with AI*'¹⁸ to give organisations practical advice on explaining the processes, services, and decisions delivered or assisted by A.I. to the individuals affected by them. The guidance listed six tasks for organisations to consider, ensuring these elements are incorporated into the A.I. solution's design and deployment. To ensure that consumers, businesses, and the wider economy

reap the benefits of developments in A.I. while harms are mitigated, the Competition & Markets Authority issued in 2024 the '*CMA AI Strategic Update*'¹⁹, mandating developers and deployers to take responsibility for their parts of the value chain.

- **US:** The NIST's '*Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*' suggests that governance tools and protocols for other A.I. systems can also be utilised for GenA.I., including accessibility measures, A.I. actor qualifications, auditing, change management, incident response, risk controls, and monitoring. Furthermore, organisations should adopt explainability methods to assess GenA.I. solutions' outputs and ensure alignment with their intended purposes.

Governance and accountability are critical components of any regulatory framework for GenA.I. systems. By delineating responsibility and offering explainability considerations, jurisdictions are working to promote governance and accountability to mitigate risks associated with A.I. technologies. These measures not only enhance the reliability and safety of A.I. systems but also foster public trust and confidence in the responsible use of A.I., facilitating more informed decision-making in A.I. deployment.

4.2. Fairness

While A.I. has the potential to address specific problems, it does not possess an inherent understanding of morality, a distinctly human attribute. Without careful oversight from developers and service providers, particularly in filtering training data and establishing clear guidelines for prompts and outputs, A.I. systems may produce biased or discriminatory content. Such outcomes can negatively impact users and potentially undermine public trust in the technology.

¹⁷ Infocomm Media Development Authority and Personal Data Protection Commission. 2020. *Model Artificial Intelligence Governance Framework Second Edition*. (<https://www.pdpc.gov.sg/-/media/files/pdpc/pdf-files/resource-for-organisation/ai/smodelaigovframework2.pdf>).

¹⁸ Information Commissioner's Office and The Alan Turing Institute. 2022. *Explaining Decisions Made with AI*. (<https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/explaining-decisions-made-with-artificial-intelligence/>).

¹⁹ Competition & Markets Authority. 2024. *CMA AI Strategic Update*. (<https://www.gov.uk/government/publications/cma-ai-strategic-update/cma-ai-strategic-update>).

To address these concerns, guidelines and frameworks have been issued to ensure the fairness of A.I.-generated outputs. These documents emphasise the need for solution providers to treat all users equally and inclusively and avoid discrimination or bias (see Figure 23).

Figure 23: Fairness

OBJECTIVE

Ensure equal and inclusive treatment of all users, prevent bias, and avoid discrimination against any certain group

JURISDICTION	RELEVANT EXCERPTS
European Union High-Level Expert Group on Artificial Intelligence • <i>Ethics guidelines for trustworthy AI (2019)</i>	• Identifiable and discriminatory bias should be removed in the collection phase where possible. The way in which AI systems are developed ... may also suffer from unfair bias. This could be counteracted by putting in place oversight processes to analyse and address the system's purpose, constraints, requirements and decisions in a clear and transparent manner. • Particularly in business-to-consumer domains, systems should be user-centric and designed in a way that allows all people to use AI products or services...
Hong Kong (SAR) Digital Policy Office of the Hong Kong Government • <i>Ethical Artificial Intelligence Framework (2024)</i>	• The recommendation/result from the AI applications should treat individuals within similar groups in a fair manner, without favouritism or discrimination and without causing or resulting in harm. • Organisations can mitigate inherent bias using bias detection methods to evaluate whether the model discriminates against various groups or individuals, given a test-dataset.
Mainland China National Technical Committee 260 on Cybersecurity • <i>Basic Security Requirements for Generative Artificial Intelligence Service (2024)</i>	• Discriminatory contents are a major security risk for data and generated contents, including those by ethnicity, belief, country, region, gender, age, profession, health condition, and other aspects. • In terms of corpus content filtering: keywords, classification models, manual sampling and other methods should be adopted to fully filter out illegal and harmful information in all corpora.
Singapore Infocomm Media Development Authority and Personal Data Protection Commission • <i>Model Artificial Intelligence Governance Framework Second Edition (2020)</i>	• While identifying and addressing inherent bias in datasets may not be easy, organisations can mitigate the risk of inherent bias by having a heterogeneous dataset (i.e., collecting data from a variety of reliable sources). • Where applicable, the model could also be checked for systematic bias by testing it on different demographic groups to observe whether any groups are being systematically advantaged or disadvantaged.
United Kingdom Office for Artificial Intelligence and Department for Science, Innovation & Technology • <i>A Pro-innovation Approach to AI Regulation (2023)</i>	• AI systems should not undermine the legal rights of individuals or organisations, discriminate unfairly against individuals or create unfair market outcomes. Actors involved in all stages of the AI life cycle should consider definitions of fairness that are appropriate to a system's use, outcomes and the application of relevant law.
United States National Institute of Standards and Technology • <i>Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (2024)</i>	• Conduct fairness assessments to measure systemic bias. Measure GAI (GenAI.) system performance across demographic groups and subgroups, addressing both quality of service and any allocation of services and resources. • Quantify harms using: field testing with sub-group populations to determine likelihood of exposure to generated content exhibiting harmful bias, AI red-teaming with counterfactual and low-context (e.g., "leader," "bad guys") prompts.

Source: Regulatory disclosures, Quinlan & Associates analysis

- **EU:** The High-Level Expert Group on Artificial Intelligence²⁰ presented the '*Ethics guidelines for trustworthy AI*'²¹, which outlines seven key requirements that A.I. systems must meet to be considered trustworthy. One of these requirements emphasises the need to eliminate identifiable and discriminatory biases as much as possible during the data collection phase. Additionally, the development of A.I. systems, including the programming of algorithms, can introduce unfair biases, which necessitates the implementation of robust governance and oversight practices. Specifically, in business-to-consumer contexts, A.I. systems should be user-centric and designed to be accessible to all individuals, regardless of age, gender, abilities, or other characteristics.
- **Hong Kong (SAR):** The '*Ethical Artificial Intelligence Framework*', issued by the Digital Policy Office of the Hong Kong Government, provides guiding principles, practices, and an assessment template for the government and other organisations. The framework states that A.I. recommendations and results should


treat individuals within similar groups fairly, without favouritism or discrimination, and without causing harm. Practically, it advises organisations to use bias detection methods to evaluate whether A.I. models discriminate against specific groups or individuals using a test dataset.

- **Mainland China:** The National Technical Committee 260 on Cybersecurity has outlined the '*Basic Security Requirements for Generative Artificial Intelligence Service*'²², which provides a wide range of examples of potential discrimination and misuse. These examples, including content that discriminates based on ethnicity, beliefs, or gender, help service providers clearly understand the scope of discriminatory content. To prevent such undesirable outcomes, proactive measures should be implemented from the outset. This includes using keywords, classification models, and manual sampling to thoroughly filter illegal and undesirable information throughout the entire training dataset.

²⁰ The European Commission appointed a group of experts to form the High-level Expert Group on Artificial Intelligence to provide advice on its A.I. strategy. The concept of trustworthiness and the 7 key requirements, introduced by the '*Ethics guidelines for trustworthy AI*', are guiding the upcoming legislative steps in A.I. for the European Commission.

²¹ High-Level Expert Group on Artificial Intelligence. 2019. *Ethics guidelines for trustworthy AI*. (<https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>).

²² National Technical Committee 260 on Cybersecurity. 2024. *Basic Security Requirements for Generative Artificial Intelligence Service*. (Chinese - <https://www.tc260.org.cn/upload/2024-03-01/1709282398070082466.pdf>).



- Singapore:** Singapore's *'Model Artificial Intelligence Governance Framework Second Edition'* provides comprehensive and actionable guidance for private sector organisations to navigate the ethical and governance challenges associated with deploying A.I. solutions. One of the critical challenges addressed by the framework is the identification and mitigation of inherent bias in datasets. Although detecting such biases can be complex, organisations can reduce this risk by ensuring their datasets are diverse and incorporating data from a broad range of reliable sources. In addition, whenever relevant, A.I. models should be systematically evaluated for potential bias by testing their performance across different demographic groups. This allows organisations to identify whether certain groups are being consistently advantaged or disadvantaged, enabling them to take corrective measures to improve fairness and inclusivity.
- U.K.:** The command paper *'A Pro-innovation Approach to AI Regulation'*²³, issued by the Office for Artificial Intelligence and Department for Science, Innovation and Technology, emphasises that A.I. systems should not infringe on the legal rights of individuals or organisations, engage in unfair discrimination, or create unjust market outcomes. It advises all stakeholders involved in the A.I. lifecycle to consider definitions of fairness that align with the system's intended use, outcomes, and applicable laws.
- US:** The *'Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile'* issued by the NIST recommends conducting fairness assessments to identify systemic bias in A.I. systems. It advises measuring the performance of GenA.I. systems across various demographic groups and subgroups, considering both the quality of service and the distribution of services and resources. The framework also suggests quantifying potential harms through field testing

with subgroup populations to assess exposure to biased content and employing A.I. red-teaming techniques using counterfactual and low-context prompts (such as "leader" or "bad guys").

Fairness is a critical principle in the development and deployment of GenA.I., ensuring that these technologies do not perpetuate bias or discrimination and are used responsibly. By issuing clear guidelines on the ethical use of A.I., regulators across different jurisdictions are working to mitigate risks associated with bias and discrimination. These measures are essential not only for protecting users but also for maintaining public trust in A.I. technologies.

4.3. Data Privacy & Protection

The increasing sophistication in GenA.I. systems has escalated the demand for vast volumes of high-quality data to train these algorithms. This surge in data usage raised critical concerns related to privacy, security, and the responsible stewardship of sensitive information. Striking the right balance between the benefits of data-driven innovation and the imperative to protect data security, individual privacy, and copyright remains a pressing challenge.

Relevant authorities have mandated the establishment of protections for the legitimate collection and handling of data in A.I. and emphasised the importance of adopting cybersecurity measures and adhering to current regulations. Ensuring data privacy and protection builds trust with users, safeguards personal information from unauthorised access, and enhances the ethical use of A.I., contributing to responsible innovation (see Figure 24).

²³ Office for Artificial Intelligence and Department for Science, Innovation & Technology. 2023. *A Pro-innovation Approach to AI Regulation*. (<https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper#correction-slip>).

Figure 24: Data Privacy & Protection

OBJECTIVE

Implement safety and cybersecurity measures for lawful data collection to ensure data privacy and security

JURISDICTION	RELEVANT EXCERPTS
European Union European Union Agency for Cybersecurity • <i>Multilayer Framework for Good Cybersecurity Practices for AI</i> (2023) European Data Protection Board ChatGPT Taskforce • <i>Report of the Work Undertaken by the ChatGPT Taskforce</i> (2024)	• The main recommendation is to treat the cybersecurity of AI systems as an additional effort to existing practices, ... with AI-specific practices. • ... defining precise collection criteria and ensuring that certain data categories are not collected or that certain sources are excluded from data collection ... delete or anonymise personal data ... before the training stage... Data subjects should ... be clearly and demonstrably informed that such "Content" may be used for training purposes.
Hong Kong (SAR) Office of the Privacy Commissioner for Personal Data • <i>Guidance on the Ethical Development and Use of Artificial Intelligence</i> (2021) • <i>Artificial Intelligence: Model Personal Data Protection Framework</i> (2024)	• Personal data involved in the development and use of AI should be processed and protected in accordance with the PDPO (Personal Data (Privacy) Ordinance)... and cover the entire life cycle of the handling of personal data... • AI systems should be monitored and reviewed continuously because the risk factors related to their use may change over time... High-risk AI necessitate more frequent and stringent monitoring and review... Organisations should consider incorporating the following review mechanisms...
Mainland China Cyberspace Administration of China & 6 other Ministries ¹ • <i>Interim Measures for the Administration of Generative Artificial Intelligence Services</i> (2023) National Technical Committee 260 on Cybersecurity • <i>Basic Security Requirements for Generative Artificial Intelligence Service</i> (2024)	• The consent of the individual shall be obtained if personal information is involved. The provider shall fulfill its obligation to protect the user's input information and usage records lawfully, ..., shall not illegally retain the input information and usage records that can identify the user's identity, and shall not illegally provide the user's input information and usage records to others. • Regular security audits should be conducted on the development frameworks and codes used ... Potential security holes should be identified and repaired.
Singapore Personal Data Protection Commission • <i>Advisory Guidelines on Use of Personal Data in AI Recommendation and Decision Systems</i> (2024) Infocomm Media Development Authority • <i>Model AI Governance Framework for Generative AI: Fostering a Trusted Ecosystem</i> (2024)	• ... appropriate technical, process and / or legal controls for data protection should be included. Where possible, organisations are encouraged to pseudonymise or de-identify the personal data used as a basic control. • As required under their Protection Obligation, Service Providers should guard against unauthorised modification of the personal data they are processing. • Security-by-design is a fundamental security concept... New tools have to be developed ... to help enhance the ability to identify and extract malicious codes...
United Kingdom Information Commissioner's Office • <i>Generative AI: Eight Questions that Developers and Users Need to Ask</i> (2023) National Cyber Security Centre • <i>Guidelines for Secure AI System Development</i> (2023)	• Organisations developing or using generative AI should be considering their data protection obligations from the outset, taking a data protection by design and by default approach... Data protection law still applies when the personal information that you're processing comes from publicly accessible sources. • Security must be a core requirement, not just in the development phase, but throughout the life cycle of the system.
United States National Institute of Standards and Technology • <i>Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile</i> (2024)	• Conduct periodic monitoring of AI-generated content for privacy risks; address any possible instances of PII or sensitive data exposure... Leverage approaches to detect the presence of PII or sensitive data in generated output text, image, video, or audio. • Establish policies to evaluate risk-relevant capabilities of GAI and robustness of safety measures, both prior to deployment and on an ongoing basis, through internal and external evaluations.

¹The six ministries are: National Development and Reform Commission, Ministry of Education, Ministry of Science and Technology, Ministry of Industry and Information Technology, Ministry of Public Security, and National Radio and Television Administration.

Source: Regulatory disclosures, Quinlan & Associates analysis

- EU:** The European Union Agency for Cybersecurity has developed a '*Multilayer Framework for Good Cybersecurity Practices for AI*'²⁴, which recommends stakeholders enhance existing protocols by integrating dynamic risk assessments throughout the A.I. lifecycle. In response to the rise of generative A.I., the European Data Protection Board formed a ChatGPT Taskforce, which issued '*Report of the Work Undertaken by the ChatGPT Taskforce*'²⁵ in May 2024. The report highlighted that data controllers for personal data related to large language models must comply with regulations, establish data collection criteria, anonymise data from web scraping, and inform users about the potential use of their input for training.
- Hong Kong (SAR):** In response to the rapid development and adoption of A.I. solutions, the Office of the Privacy Commissioner for Personal Data issued the '*Guidance on the Ethical Development and Use of Artificial Intelligence*'²⁶, which aims to provide organisations with recommendations to embrace key data stewardship and ethical values. It stated privacy is a fundamental human right and requires compliant and effective data governance to protect individuals' privacy. The Guidance also reminded A.I. suppliers of the need to fulfil user requests regarding data access and correction. Building upon the 2021 Guidance, the 2024 '*Artificial Intelligence: Model Personal Data Protection Framework*'²⁷ was published to provide practical recommendations with a risk-based approach to ensure compliance. It also highlights the need to pay special attention to high-risk A.I. systems and suggests review mechanisms such as documentation, monitoring, re-assessments, and periodic reviews.
- Mainland China:** To ensure the healthy development and regulated use of GenA.I., 7 Chinese ministries issued the '*Interim Measures for the Administration of Generative Artificial Intelligence Services*'²⁸ in accordance with relevant laws, protecting national security, public interests, and the rights of individuals and organisations. Consent from individuals is required for handling personal information. The provider must legally protect user input and usage records, refrain from unlawfully retaining identifiable information, and must not share this information with unauthorised parties. In addition, the '*Basic Security Requirements for Generative Artificial Intelligence Service*' listed the need for regular security audits, focusing on open-source frameworks' vulnerabilities and identifying potential security holes.
- Singapore:** Following a public consultation, Singapore's Personal Data Protection Commission issued the '*Advisory Guidelines on Use of Personal Data in AI Recommendation and Decision Systems*'²⁹, which suggested that organisations should include appropriate technical, process and/or legal controls for data protection when designing, training, testing, or monitoring A.I. systems using personal data. Organisations are encouraged to pseudonymise or de-identify personal data as a fundamental measure. Service providers must also protect against unauthorised modifications of the personal data they process, in line with their protection obligations. Moreover, the IMDA's '*Model AI Governance Framework for Generative AI: Fostering a Trusted Ecosystem*'³⁰ also highlighted the need for new tools to help enhance cybersecurity, including input filters and digital forensics tools.

²⁴ European Union Agency for Cybersecurity. 2023. *Multilayer Framework for Good Cybersecurity Practices for AI*. (<https://www.cybersecitalia.it/wp-content/uploads/2023/06/Multilayer-Framework-for-Good-Cybersecurity-Practices-for-AI.pdf>).

²⁵ European Data Protection Board ChatGPT Taskforce. 2024. *Report of the Work Undertaken by the ChatGPT Taskforce*. (https://www.edpb.europa.eu/system/files/2024-05/edpb_20240523_report_chatgpt_taskforce_en.pdf).

²⁶ Office of the Privacy Commissioner for Personal Data. 2021. *Guidance on the Ethical Development and Use of Artificial Intelligence*. (https://www.pcpd.org.hk/english/resources_centre/publications/files/guidance_ethical_e.pdf).

²⁷ Office of the Privacy Commissioner for Personal Data. 2024. *Artificial Intelligence: Model Personal Data Protection Framework*. (https://www.pcpd.org.hk/english/resources_centre/publications/files/ai_protection_framework.pdf).

²⁸ Cyberspace Administration of China, the National Development and Reform Commission, the Ministry of Education, the Ministry of Science and Technology, the Ministry of Industry and Information Technology, the Ministry of Public Security, and the National Radio and Television Administration. 2023. *Interim Measures for the Administration of Generative Artificial Intelligence Services*. (https://www.cac.gov.cn/2023-07/13/c_1690898327029107.htm).

²⁹ Personal Data Protection Commission. 2024. *Advisory Guidelines on Use of Personal Data in AI Recommendation and Decision Systems*. (<https://www.pdpc.gov.sg/-/media/files/pdpc/pdf-files/advisory-guidelines/advisory-guidelines-on-the-use-of-personal-data-in-ai-recommendation-and-decision-systems.pdf>).

³⁰ Infocomm Media Development Authority. 2024. *Model AI Governance Framework for Generative AI: Fostering a Trusted Ecosystem*. (<https://aiverifyfoundation.sg/wp-content/uploads/2024/05/Model-AI-Governance-Framework-for-Generative-AI-May-2024-1-1.pdf>).

- **U.K.:** The Information Commissioner's Office has published '*Generative AI: eight questions that developers and users need to ask*'³¹ to help developers and users navigate relevant supervisory requirements. Stakeholders must prioritise data privacy by adopting a "data protection by design and by default" approach to consider obligations from the very beginning of their projects. Importantly, data protection laws still apply even when personal information is obtained from publicly accessible sources. In addition, the National Cyber Security Centre has issued the '*Guidelines for Secure AI System Development*'³² for all stakeholders to make informed decisions about the design, development, deployment, and operation of their A.I. systems. The guidelines emphasise a "secure by default" approach, prioritising customer security and leadership commitment to security throughout the lifecycle.
- **US:** Developed in response to legislative mandates, the '*Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*' establishes a comprehensive, structured approach to identifying and mitigating risks associated with A.I. technologies. The Profile highlighted the need for periodic monitoring of privacy risks

and methods to identify personally identifiable information in generated outputs. The framework also suggests risk evaluation prior to deployment and on an ongoing basis to ensure cybersecurity.

As the deployment of GenA.I. continues to expand, data privacy and protection are becoming increasingly critical. Responsible and secure data handling is essential for building trust, maintaining regulatory compliance, and safeguarding the integrity of the financial services industry as it integrates GenA.I. systems into its operations.

One of the key findings from our interviews is that financial institutions are actively seeking greater clarity on the responsible use of sensitive data in refining the outputs generated by GenA.I. solutions. During the interviews, several financial institutions indicated that clear regulatory guidance on the use of sensitive data, including data classification standards and the extent to which they can be leveraged to refine GenA.I. outputs, would greatly enhance the utilisation of GenA.I. capabilities. These institutions are particularly concerned with ensuring that these outputs remain both relevant and up-to-date, enabling improved service delivery to customers, while maintaining strict adherence to regulatory compliance.

³¹ Information Commissioner's Office. 2023. *Generative AI: Eight Questions that Developers and Users Need to Ask*. (<https://ico.org.uk/about-the-ico/media-centre/blog-generative-ai-eight-questions-that-developers-and-users-need-to-ask/>).

³² National Cyber Security Centre. 2023. *Guidelines for Secure AI System Development*. (<https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development>).



4.4. Transparency & Disclosure

With the increasing use of A.I. in the financial services industry, customers may not always be fully aware of the extent to which A.I. systems are involved in service delivery. This lack of transparency can

affect customer decision-making, where trust and informed consent are essential. Financial regulators are recognising the need for institutions to disclose their use of A.I. more explicitly, ensuring customers understand how A.I. is integrated into services to maintain trust and safeguard consumer rights (see Figure 25).

Figure 25: Transparency & Disclosure

OBJECTIVE

Provide users with clear and sufficient disclosures about A.I. usage in the services to protect user rights

JURISDICTION	RELEVANT EXCERPTS
European Union European Securities and Markets Authority Public Statement on the Use of Artificial Intelligence in the Provision of Retail Investment Services (2024)	- Furthermore, investment firms should be transparent on the role of AI in investment decision-making processes related to the provision of investment services. - In general, the ESMA expects that when firms provide clients with information on how the firms use AI tools for the provision of investment services, they ensure that such information is presented in a clear, fair and not misleading manner.
Hong Kong (SAR) Hong Kong Monetary Authority - High-level Principles on Artificial Intelligence (2019) - Consumer Protection in Respect of Use of Generative Artificial Intelligence (2024) Insurance Authority - Conduct in Focus - Chatting about Chatbots and AI (2023)	- To provide transparency and thereby increase consumers' confidence in A.I.-powered services, banks should make clear to the consumer, prior to service provision, that the relevant service is powered by A.I. technology and the risks involved. - Accordingly, they (i.e., banks) should disclose the use of GenA.I. to customers, and, among others, communicate with customers on the use and purpose of adoption of the GenA.I. models as well as the limitations of such models, in order to enhance customers' understanding of the model-generated outputs. Clear disclosure would need to be made as to the Chatbot's limitations, how it should be used, the data set it is trained on and how that data is stored, used and how long it is kept.
Mainland China People's Bank of China Guidance on Information Disclosure for Financial Applications Based on Artificial Intelligence Algorithms (2023)	Artificial intelligence algorithm financial application disclosure content is divided into seven aspects: algorithm combination disclosure, algorithm logic disclosure, algorithm application disclosure, algorithm data disclosure, algorithm subject disclosure, algorithm change disclosure, and algorithm audit disclosure.
Singapore Monetary Authority of Singapore Emerging Risks and Opportunities of Generative AI for Banks A Singapore Perspective (2023)	- Transparency principles require disclosing the use of Artificial Intelligence and Data Analytics, explaining outcomes, and understanding what data is used to drive the outcome. - Data subjects are provided, upon request, clear explanations of the on sequences that AIDA (Artificial Intelligence and Data Analytics)-driven decisions may have on them.
United Kingdom Financial Conduct Authority AI Update (2024)	- Whilst our regulatory framework does not specifically address the transparency or explainability of AI systems, there are a number of high-level requirements and principles under our approach to consumer protection. - Related rules under the Consumer Duty on consumer understanding refer to meeting the information needs of retail customers and equipping them to make decisions that are effective, timely and properly informed.
United States National Institute of Standards and Technology Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (2024)	Consider disclosing use of GAI (i.e., GenA.I.) to end users in relevant contexts, while considering the objective of disclosure, the context of use, the likelihood and magnitude of the risk posed, the audience of the disclosure, as well as the frequency of the disclosures.

Source: Regulatory disclosures, Quinlan & Associates analysis

- EU:** In its *'Public Statement on the Use of Artificial Intelligence in the Provision of Retail Investment Services'*³³, the European Securities and Markets Authority (ESMA) highlights both the opportunities and risks associated with A.I. in the delivery of investment services. A key directive from ESMA is that investment firms should be transparent about the role A.I. plays in their decision-making processes. Firms are expected to communicate their use of A.I. tools in a manner that is clear, fair, and not misleading, ensuring that clients are fully informed and can make well-founded decisions regarding their investments.
- Hong Kong (SAR):** In 2019, the HKMA issued a circular, titled *'High-level Principles on Artificial Intelligence'*³⁴, to provide guidance on the use of A.I. in banking sector. The circular underscores the need for transparency and disclosure, stating that customers must be clearly informed when relevant services utilise A.I. technology and any associated risks. Recently, the HKMA issued another circular, *'The Consumer Protection in respect of Use of Generative Artificial Intelligence'*³⁵, providing guidance on using GenA.I. in customer-facing applications. The circular reinforces the importance of transparency, requiring authorised institutions to provide clear, accurate, and easily understandable disclosures about their use of GenA.I. solutions. This includes informing customers about the purpose and limitations of the A.I. models, helping them better understand the generated outputs. For the insurance sector, the IA's Circular on *'Conduct in Focus – Chatting about Chatbots and AI'*³⁶ discusses key considerations and perspectives on A.I.-powered Chatbots under the insurance regulated activities regime. The paper highlights the importance of clear disclosure to ensure customers are not only aware of the Chatbot's limitations and data management processes, but also able to make fully informed decisions in their best interests.
- Mainland China:** The *'Guidance on Information Disclosure for Financial Applications Based on Artificial Intelligence Algorithms'*³⁷, issued by the People's Bank of China, outlines key principles, formats, and content for disclosing the use of A.I. algorithms in the financial sector. The document aims to better protect the legitimate rights and interests of financial consumers by ensuring transparency in A.I.-driven financial services. It identifies seven critical areas for disclosure: algorithm combination, algorithm logic, algorithm application, algorithm data, algorithm subject, algorithm changes, and algorithm auditing, specifying the items to be disclosed under each area. Additionally, the guidance includes a template that financial institutions can reference to ensure proper disclosure of A.I. algorithm-related information.
- Singapore:** The Monetary Authority of Singapore has released a whitepaper titled *'Emerging Risks and Opportunities of Generative AI for Banks, A Singapore Perspective'*³⁸ as part of Project MindForge. In the first phase of the project, a comprehensive framework for managing GenA.I. risks was developed. The whitepaper emphasises that transparency principles require disclosing the use of Artificial Intelligence and Data Analytics (AIDA), explaining the outcomes produced, and clarifying the data driving these outcomes. Additionally, data subjects must be provided with clear explanations, upon request, about the potential impacts of AIDA-driven decisions on them.
- U.K.:** Following the Government's publication of its pro-innovation strategy on A.I., the Financial Conduct Authority (FCA) outlined its approach to A.I. in its *'AI Update'*³⁹. While the FCA acknowledges that its regulatory framework does not specifically address the transparency or explainability of A.I. systems, it upholds several high-level requirements and principles under its consumer protection approach that are pertinent

³³ European Securities and Markets Authority. 2024. Public Statement on the Use of Artificial Intelligence in the Provision of Retail Investment Services. (https://www.esma.europa.eu/sites/default/files/2024-05/ESMA35-335435667-5924_Public_Statement_on_AI_and_investment_services.pdf).

³⁴ Hong Kong Monetary Authority. 2024. High-level Principles on Artificial Intelligence. (<https://www.hkma.gov.hk/media/eng/doc/key-information/guidelines-and-circular/2019/20191101e1.pdf>).

³⁵ Hong Kong Monetary Authority. 2024. Consumer Protection in Respect of Use of Generative Artificial Intelligence. (<https://www.hkma.gov.hk/media/eng/doc/key-information/guidelines-and-circular/2024/20240819e1.pdf>).

³⁶ Insurance Authority. 2023. Conduct in Focus - Chatting about Chatbots and AI. (https://www.ia.org.hk/en/legislative_framework/circulars/reg_matters/files/Cir_dd_12.05.2023_Eng.pdf).

³⁷ People's Bank of China. 2023. Guidance on Information Disclosure for Financial Applications Based on Artificial Intelligence Algorithms. (<http://www.hbhill.com/uploadFiles/-13/395/070/09/2c9081e489c0507d018cdeab565401fc.pdf>).

³⁸ Monetary Authority of Singapore. 2023. Emerging Risks and Opportunities of Generative AI for Banks, A Singapore Perspective. (<https://www.mas.gov.sg/-/media/mas-media-library/schemes-and-initiatives/ftig/project-mindforge/emerging-risks-and-opportunities-of-generative-ai-for-banks.pdf>).

³⁹ Financial Conduct Authority. 2024. AI Update. (<https://www.fca.org.uk/publication/corporate/ai-update.pdf>).

to the information provided to consumers. These principles are relevant for firms using A.I. responsibly and safely in financial services. Notably, the Consumer Duty imposes a broad obligation to act in good faith, characterised by honesty and fair, open dealings with retail consumers. Additionally, the Consumer Duty includes rules on consumer understanding, which focus on meeting the information needs of retail customers and enabling them to make effective, timely, and well-informed decisions.

- **US:** The '*Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*' issued by the NIST recommends that organisations consider disclosing the use of GenA.I. to end-users in appropriate contexts. This disclosure is suggested to be thoughtfully tailored, taking into account several key factors: the purpose and objective of the disclosure, the specific context in which the GenA.I. is being used, the likelihood and severity of any associated risks, the characteristics of the intended audience, and the optimal frequency of such disclosures. These considerations help ensure that users are informed in a transparent and meaningful way without causing unnecessary confusion or information overload.

As transparency and disclosure principles evolve across various jurisdictions, solution providers and financial institutions must adjust to meet new regulatory expectations. These regulations are designed to protect user rights and promote responsible A.I. usage by ensuring that users are well-informed about how A.I. systems are used.

4.5. Conclusion

It is imperative for financial institutions to remain informed about these regulatory updates, as they provide guidance on suggested practices required for ongoing innovation efforts, particularly for those operating on a regional or global scale. Keeping updated on these regulations ensures that organisations can effectively adapt and comply, maintaining their competitive edge in an increasingly complex regulatory environment.


Staying ahead of the curve is equally critical for the public sector. By continuously revisiting and updating existing regulatory frameworks and incorporating global suggested practices, where appropriate, the public sector can establish effective policies that not only drive innovation but also protect the interests of the broader ecosystem.

5. Adoption Considerations and Suggested Practices

Given that GenA.I. is a relatively novel technology, financial institutions have expressed the need for further guidance on how to effectively plan, test, deploy, and integrate these solutions with appropriate data in their business operations. This section aims to address these concerns by offering a comprehensive list of key adoption considerations and suggested practices across the deployment value chain, specifically tailored to GenA.I. solutions for financial institutions.

First and foremost, a robust GenA.I. governance structure must be established to provide the overarching guidance necessary for compliant and sustainable deployment initiatives. Once a robust governance structure is in place, financial institutions can look to institute suggested practices across the deployment value chain, namely: (1) pre-deployment, (2) deployment, and (3) post-deployment. Finally, a strong feedback loop should be implemented across all phases to support the collection, analysis, and implementation of feedback (see Figure 26).

Figure 26: GenA.I. Adoption Considerations




Source: Quinlan & Associates analysis

5.1. GenA.I. Governance Structure

Establishing a robust governance structure is essential for the successful deployment of GenA.I. solutions at any financial institution, ensuring that the technology is used both correctly and responsibly.

Taking into consideration the regulatory standards and supervisory guidelines outlined across major jurisdictions in Section 4, including the recent circular on GenA.I. published by the HKMA⁴⁰, a range of key measures should be adopted by financial institutions to establish robust GenA.I. governance (see Figure 27).

Figure 27: GenA.I. Governance Structure



Source: Regulatory disclosures, Quinlan & Associates analysis

- Governance & Accountability:** Financial institutions should establish a robust governance, oversight, and accountability framework to use GenA.I. solutions. The framework should define clear roles and responsibilities across the development lifecycle, guide thorough risk assessments with ongoing monitoring, and put in place rigorous incident tracking and response mechanisms. The board and senior management ultimately remain accountable for decisions made using

GenA.I., which must comply with consumer protection principles, corporate values, and ethical standards. Moreover, financial institutions need to ensure an appropriate level of explainability of the GenA.I. solutions and validate the solutions before launch. Appropriate policies and procedures should also be in place to guarantee responsible use. During the early stage of deploying GenA.I. applications, "human-in-the-loop" approach should also be adopted (i.e., having human to retain control in the decision-making process).

⁴⁰ Hong Kong Monetary Authority. 2024. *Consumer Protection in respect of Use of Generative Artificial Intelligence*. (<https://www.hkma.gov.hk/media/eng/doc/key-information/guidelines-and-circular/2024/20240819e1.pdf>).

- **Fairness:** GenA.I. models must produce objective, consistent, ethical, and fair outcomes. Financial institutions should comply with applicable laws, consider customers' financial capabilities and needs, and appropriately weigh all relevant variables to build a robust model. Unfair bias or disadvantage must be avoided through techniques like anonymisation, fair representation, and bias adjustments. Customers should be given the option to opt-out and request human intervention during the early stage of deploying GenA.I. applications. Rigorous adversarial testing should be conducted to identify vulnerabilities and protect against potential manipulation.
- **Data Privacy & Protection:** Effective data protection measures must be implemented to safeguard customer information. Financial institutions should embed data protection in the product and system design. Consent should be obtained with a clear understanding before the collection and use of personal data. Continuous monitoring and auditing of A.I. systems are crucial to ensure cybersecurity, ongoing compliance, and system integrity, with rigorous data governance practices and secure-by-design principles to minimise vulnerabilities and protect against malicious attacks.
- **Transparency & Disclosure:** Financial institutions should provide an appropriate level of transparency around the adoption, nature, usage, risks, and limitations of the GenA.I. solution. End-users should be able to enquire about and request reviews of the system. They should also receive adequate support with accessible and fair complaint handling. GenA.I. service providers should explain the types of data used and factors behind A.I.-driven decisions, enhancing consumer understanding through appropriate education. Communications should be clear and simple to understand, enabling customers to interpret the outputs and build trust.

A well-defined governance structure is essential for ensuring clear oversight and accountability throughout GenA.I. initiatives. Such structure must explicitly delineate the roles and responsibilities of all stakeholders, including project managers, Information Technology (IT) teams, and business leaders. It should also outline mechanisms for effective decision-making, comprehensive risk management, and efficient conflict resolution.


Several financial institutions we interviewed have already implemented GenA.I. governance committees, uniting key stakeholders to thoroughly review GenA.I. initiatives. These committees provide early insight into innovations and facilitate feedback on relevant risks and recommended actions. One institution underscored the importance of backtesting GenA.I. solutions for both internal and external GenA.I. models before deployment, while another global bank highlighted the need for automated, continuous monitoring throughout the solution's lifecycle.

In addition to internal governance, several institutions stressed the importance of a collaborative governance structure with technology providers. Technology providers echoed this sentiment, emphasising their risk-based approach to GenA.I. governance and their commitment to working with clients to establish a clear governance framework to collectively oversee GenA.I. initiatives.

5.2. Deployment Value Chain

Anchored by a robust governance framework from the outset, organisations can start their GenA.I. deployment process that is aligned with the organisation's broader strategic objectives, helping to keep the initiative on track and within scope, ultimately facilitating adoption of the technology across the value chain: (1) pre-deployment, (2) deployment, and (3) post-deployment (see Figure 28).

Figure 28: Adoption Considerations Across the Value Chain



Source: Quinlan & Associates analysis

Throughout the entire lifecycle of a GenA.I. system, organisations should focus on two key areas: (1) management of all aspects of GenA.I. system development and (2) the implementation of processes to identify, evaluate, and monitor associated risks. The ISO / IEC 42001 International Standard for A.I. management systems⁴¹, introduced in 2023, emphasises the importance of these focus areas to ensure that systems are not only functional but also aligned with business goals and regulatory requirements. Adhering to industry-recognised standards such as the ISO / IEC 42001:2023 can foster greater trust in GenA.I. services among stakeholders.

5.2.1. Pre-deployment Phase

To initiate the deployment of a GenA.I. solution, financial institutions should first outline the rationale for undertaking the project by developing a clear business case. This should be followed by testing the concept as a Proof-of-Concept (PoC) to validate the feasibility and effectiveness of the proposed solution in a real-world context. These pre-deployment steps can help financial institutions make more informed decisions about adopting new GenA.I. solutions (see Figure 29).

Figure 29: Pre-deployment Phase



Source: Quinlan & Associates analysis

⁴¹ International Organization for Standardization. 2023. *ISO/IEC 42001:2023 Information technology — Artificial intelligence — Management system*. (<https://www.iso.org/standard/81230.html>).

Business Case Development

Identifying the problem is the first step in developing a business case. Financial institutions need to understand the nature of the issue specific to their business operations and its significance relative to industry and market metrics. A quantitative statement that clearly relates to the top and bottom lines will help define a measurable, trackable, and time-sensitive problem statement, laying the foundation for both business and technology goals. Establishing specific key performance indicators is also recommended as a detailed point of reference.

With a clear and quantified objective, financial institutions can further investigate the underlying causes of the issue. This may involve examining both macro factors (e.g., industry trends, headwinds, and tailwinds) and market factors (e.g., changes in market share and the forces affecting demand and supply). The findings should ultimately help determine whether the problem is actionable and worth addressing financially.

Building a business case for GenA.I. often starts by identifying specific tasks that require significant time commitments but have limited outcome impact, such as processing or reviewing large volumes of information to generate insights. Compared to other technology solutions on the market, both the problem statement and benefits of the solution output for GenA.I. solutions are generally intuitive to be quantified and contextualised.

Proof-of-Concept

Once the business case is successfully built, financial institutions can initiate the PoC process to ideate, validate, and develop a prototype. The PoC stage is essential as it mitigates development risks and informs future iterations, laying the foundation for successful deployment.

As the first step of the PoC, financial institutions should set the direction of the GenA.I. concept by considering the following three key elements: (1) user interface / user experience, (2) model design, and (3) capability acquisition approach.

- **User Interface (UI) / User Experience (UX):** Based on the problem statement defined during the development of the business case, the UI should be designed with end-users in mind, prioritising accessibility and intuitiveness. As outlined by the Open Web Application Security Project (OWASP)⁴², financial institutions can design user-friendly interfaces that incorporate warnings about potential inaccuracies and clearly labelled generated content, allowing users to make informed decisions about the reliability of the data. Feedback from potential end-users should be actively gathered and integrated into the solution to refine and enhance the overall experience. A user-centric approach during the PoC phase ensures higher adoption rates and user satisfaction when the solution is fully deployed at scale.
- **Model Design:** There are two types of GenA.I. models that financial institutions may consider: (1) pre-trained models and (2) customised models. The former offers a faster deployment option and can be cost-effective but may not fully address distinct business needs or provide a unique competitive edge. The latter, while more resource-intensive and time-consuming to develop, can be tailored to meet unique challenges and deliver more precise outcomes. The decision on which model to use should consider the organisation's specific use cases, available expertise, and desired level of customisation. While pre-trained, many GenA.I. models allow adjustment to a range of parameters, such as temperature, sampling methods and context-window size. Depending on the use cases, financial institutions can explore different combinations of parameters to better control outputs as well as resource consumption. Based on our interview findings, many financial institutions choose pre-trained models augmented with their internal knowledge base, achieving a balance between rapid deployment and tailored functionality for specific needs through augmentation or adaptation approaches such as RAG or Low-Rank Adaptation (LoRA). Other key considerations from the interviews include whether the GenA.I. model supports multiple languages, given that most financial institutions operate on a regional or global scale.

⁴² Open Web Application Security Project. 2023. *OWASP Top 10 for LLM Applications*. (https://owasp.org/www-project-top-10-for-large-language-model-applications/assets/PDF/OWASP-Top-10-for-LLMs-2023-slides-v1_1.pdf).

- **Capability Acquisition Approach:** Financial institutions should decide between developing the GenA.I. solution in-house or partnering with an external technology vendor. Developing the solution in-house offers greater control over the development process and allows for customisation to meet specific business needs. However, this approach demands substantial resources and expertise, as training GenA.I. models can be time-consuming and resource-intensive. In contrast, partnering with an external vendor can provide access to specialised skills and expedite the development process, though it may limit customisation options and create vendor dependency. Most financial institutions that have adopted GenA.I. solutions choose to partner with mature technology providers to leverage their expertise and reduce research and development costs. In these vendor arrangements, financial institutions should extend the controls to supplier management. This includes conducting proper due diligence to ensure that the vendor's practices align with the institution's organisational standards and setting clear GenA.I. guidelines for them, as outlined in the ISO / IEC 42001:2023.

Once the direction is set, then financial institutions should investigate business risks to navigate the inherent complexities of innovation while maintaining business continuity and safeguarding their long-term success. A proactive risk management approach assesses potential business risks and develops corresponding contingency plans.

- **Business Risk Assessment:** The adoption of new technology often comes with risks, especially in the case of GenA.I., due to its innovative nature. OWASP has identified several vulnerabilities in its Top 10 for LLM applications framework, including insecure output handling and over-reliance. In the context of GenA.I. deployment, these risks stem from an over-dependence on A.I. generated outputs without sufficient validation or oversight, leading to unintended data leaks or flawed decision-making. This is particularly critical in client-facing applications, where errors or failures could damage the trust and reputation of an organisation. From our interviews, financial institutions have already implemented internal

procedures to prevent over-reliance on GenA.I. and have developed appropriate mitigation strategies. These include incorporating a thorough human review process before any GenA.I. outputs are used and continuously monitoring system logs to ensure accuracy.

- **Business Contingency Planning:** Contingency planning during the PoC stage involves preparing for potential setbacks or failures that could occur during deployment. This includes developing backup plans for critical business functions and alternative solutions for when issues arise. Effective contingency planning ensures that the financial institution can quickly respond to unexpected challenges, maintaining business continuity and minimising the impact on operations. With sufficient preparation, financial institutions can then conduct a feasibility check before carrying out the PoC, an important step in identifying potential roadblocks upfront and determining if the project is truly worth pursuing. To evaluate whether the proposed solution has garnered enough resources and time, financial institutions can review resource requirements and lay out a deployment timeline.
- **Resource Requirements:** It is important for financial institutions to identify the necessary (1) human resources, (2) supporting hardware, and (3) capital required to execute GenA.I. initiatives. GenA.I. solutions typically consume more energy and computing power than other technology solutions, making precise resource planning even more critical. Furthermore, financial institutions should consider the resource implications of maintaining traditional or opt-out channels that some customers may prefer even after adopting GenA.I. solutions. Supporting both GenA.I. and legacy systems may require additional resources, potentially escalating costs and causing unexpected delays. Effective resource planning helps identify these challenges early, ensuring initiatives are delivered within the allocated resources and aligned with expectations.
- **Deployment Timeline:** Establishing a realistic deployment timeline is critical for managing expectations and ensuring that the PoC stays on schedule. The timeline should outline key

milestones, deadlines, and deliverables, with clear dependencies between tasks. It should also include buffers for potential delays and unforeseen challenges. A well-planned timeline helps keep the project on track, ensuring that the financial institution is prepared and on-time for the next stages of deployment.

Conclusion

These factors are crucial to ensuring that the solution aligns with the needs of the financial institutions, has the right technical capabilities, can be effectively integrated and maintained, and mitigates potential business risks. Thoroughly evaluating these aspects

during the business case and PoC stages allow financial institutions to make informed decisions and allocate resources efficiently to prepare for deploying GenA.I. solutions that meet their objectives and address their operational realities.

5.2.2. Deployment Phase

During the deployment stage, it is essential to address system integration and data migration plans, supported by robust cybersecurity measures. Ensuring a safe and efficient deployment process requires careful consideration of factors at both the organisational and solution levels (see Figure 30).

Figure 30: Deployment Phase



Source: Quinlan & Associates analysis


System Integration & Data Migration

The seamless connection between the GenA.I. solution and the organisation's existing IT infrastructure, software applications, and data sources is central to successful deployment. This includes integration with key programmes that support core business operations and associated data,

compatibility with current risk management systems, as well as ongoing performance measurement schemes.

Financial institutions should consider their technology architecture in mind, which is broken down into six layers: (1) application, (2) integration, (3) storage, (4) ingestion, (5) data sources, and (6) security (see Figure 31).

Figure 31: Architecture Design



Source: Palo IT, Quinlan & Associates analysis

All six layers must be carefully considered for the GenA.I. solution's overall functionality, efficiency, and security. Neglecting any of these layers can lead to compatibility issues, data integrity problems, communication breakdowns, and potential security vulnerabilities, which can undermine the success of the deployment.

- **Application Layer:** Most financial institutions have already adopted a wide range of technology applications, such as business intelligence and management information systems, alongside with GenA.I. solutions once they are adopted. On top of ensuring accessibility and intuitiveness as part of the UI / UX design during the PoC stage, financial institutions should also consider how the GenA.I. solution is integrated into the overall suite of applications.
- **Integration Layer:** Implementing appropriate middleware solutions, such as application programming interfaces (APIs) and data integration platforms, is the key to unlocking the full potential of GenA.I. solutions. APIs enable seamless communication between GenA.I. solutions and other supporting layers,

while data integration platforms ensure that data fed into GenA.I. models is secure, consistent and compatible.

- **Storage Layer:** When considering the storage layer for a GenA.I. solution, financial institutions should decide whether to host the solution on-premises or use a cloud-based storage infrastructure to accommodate the large volumes of data and processing required. Hosting the solution on-premises provides greater control over data and security but may require significant investment in hardware and maintenance and may have limitations on scalability. In contrast, a cloud-based solution offers scalability, flexibility, and lower infrastructure costs, although it may raise concerns about data privacy and compliance, particularly in the highly regulated financial services industry. Based on our interviews, most financial institutions are currently adopting GenA.I. solutions on-premises with no connection to the internet, as many use cases are still in the testing phase that requires a high degree of control. They are also monitoring for further regulatory clarity regarding the use of GenA.I. before scaling up these use cases to

ensure compliant adoption of the technology. Lastly, financial institutions should establish strict data governance policies to maintain data compatibility, as well as regulatory compliance, which will support data migration and the long-term maintenance required to ensure the storage effectively supports GenA.I. performance.

- **Ingestion Layer:** To handle high volumes of diverse data from various sources, such as market, transaction, and client data, financial institutions should implement powerful data streaming and processing systems that feed into the storage. These systems must be carefully stress-tested to determine whether they can handle the computational demands required GenA.I. solutions without disrupting ongoing operations.
- **Security Layer:** Ensuring data privacy is of paramount importance when deploying a GenA.I. solution, especially in financial institutions where sensitive client information may be involved. Data privacy measures should include encryption of data both in transit and at rest, as well as strict access controls to prevent unauthorised access. Financial institutions can also explore advanced encryption techniques, such as homomorphic encryption, which allows data to remain encrypted even while being processed by GenA.I. models, eliminating the need for decryption during computations. With regards to access, financial institutions can customise authorisations using OAuth2 and API keys, as recommended by OWASP. By establishing comprehensive protocols for data privacy, which should be guided by the GenA.I. governance structure, financial institutions can protect customer trust and avoid legal repercussions associated with inadvertent sensitive information disclosures and data breaches.

Technology Risk Assessment

When considering the architecture design for GenA.I. initiatives, financial institutions should rigorously identify technology-specific risks to develop corresponding risk management plans. A thorough risk assessment helps identify potential issues, vulnerabilities, and points of failure that could arise

during deployment and scaling. This ensures that financial institutions can create contingency plans to minimise the likelihood and impact of adverse events, thereby protecting both technological and financial investments and supporting the long-term success of the GenA.I. solution.

This process involves identifying a broad spectrum of potential risks, encompassing system incompatibilities, data security vulnerabilities, performance bottlenecks, and evolving threats specific to GenA.I. systems, such as data poisoning, prompt injections, and model jailbreaks. These risks, highlighted in frameworks like the OWASP Top 10 for LLMs, can introduce significant vulnerabilities if left unaddressed. For instance, data poisoning attacks can corrupt a model's training dataset, leading to compromised outputs, while prompt injections and jailbreaking can manipulate the model's behaviour, bypassing its safety mechanisms. Additionally, the integration of third-party content into GenA.I. systems, such as through RAG architectures, presents another layer of risk, as personal data leaks and unauthorised data access can occur from context-driven vulnerabilities.

To mitigate these risks, continuous monitoring—such as through intrusion detection systems (IDS) and real-time anomaly detection—is crucial for detecting suspicious activities before they cause harm. Furthermore, LLM red-teaming and penetration testing should be considered to proactively identify vulnerabilities in the model and supporting infrastructure. Red-teaming, particularly in an open-ended format, enables testers to explore a wide range of potential harms beyond predefined attack vectors, uncovering blind spots in the system's security posture. This approach allows for the documentation of previously unknown risks, such as hate speech, toxic outputs, hallucinations, and privacy violations, including the unintentional leakage of personally identifiable information (PII) from training data. By iterating on a list of harms identified through these tests, financial institutions can continually refine their risk management strategies and strengthen their defences.

Moreover, application-layer threats often require a different red-teaming focus, as tool-based vulnerabilities (e.g., SQL injections, privilege escalations, or hijacking) are more likely to occur once the model is embedded in operational environments.

Both white-box and black-box testing should be employed. White-box testing can uncover deep, structural weaknesses in the model, such as model poisoning or access to PII, while black-box testing is better suited to simulate real-world attack scenarios, revealing unexpected behaviours in how the model interacts with external systems.

Technology Contingency Planning

In addition to proactive risk identification, financial institutions should engage in technology contingency planning, which can be both pre-emptive and reactive. Pre-emptive measures, such as always-on firewalls, data encryption, and real-time anomaly detection systems, aim to address risks before they materialise. However, given the complexity of GenA.I. systems, reactive contingency plans must also be in place. These plans, developed in response to risk assessments, should include backup solutions, such as data recovery systems, alternative workflows, and incident response protocols. This is particularly critical in addressing application-layer threats, which often manifest after the model has been integrated into broader systems. Issues like indirect prompt injections, data exfiltration, and unauthorised access to APIs or databases may arise in real-world usage scenarios, making it essential for these plans to cover not only model failures but also broader system disruptions.

Finally, financial institutions should ensure that their contingency plans and mitigations are regularly tested through performance testing, stress testing, and load testing. This iterative process helps validate that backup systems and security measures function effectively when needed. As part of this ongoing risk management cycle, institutions should continue conducting guided red-teaming to probe for known harms, while remaining vigilant for new threats that may emerge.

Solution Performance Testing

Prior to full-scale deployment, performance testing is essential to verify that the GenA.I. solution meets operational standards. This testing evaluates the solution's ability to handle the expected volume and complexity of prompts, while maintaining accuracy, ethics, and relevance under real-world conditions. It also assesses the system's resilience to potential

disruptions or unexpected scenarios. The results guide financial institutions in validating that the solution aligns with their objectives and operational demands. If any weaknesses are found, an iterative testing process allows for targeted adjustments, ensuring the solution operates consistently and with high quality before being widely implemented.

Additionally, stress testing and load testing are critical in ensuring the solution's robustness. Stress testing pushes the system beyond its typical capacity, simulating extreme conditions like high query loads or infrastructure outages to identify breaking points and assess recovery capabilities. This helps ensure the system can operate effectively in high-pressure scenarios, protecting operational stability and the client experience. Meanwhile, load testing measures how well the system performs under expected and peak usage, assessing its ability to maintain response times and accuracy at scale, taking into consideration the considerable amount of computational power required to operate the GenA.I. solution. These tests help institutions identify potential bottlenecks and optimise resources, ensuring the GenA.I. solution can scale efficiently as demand grows.

Lastly, GenA.I. solutions should be thoroughly tested and validated under pre-defined scenarios to ensure that key risk areas identified during the technology risk assessment stage (e.g., hate speech, toxic outputs, hallucinations, and privacy violations) are adequately addressed. This approach helps identify potential risk incidents arising from GenA.I.-powered operations, enabling financial institutions to deploy these solutions safely and responsibly.

Conclusion

Throughout the solution's technical design, risk management, and performance testing, successful system integration plays a pivotal role. This involves embedding the GenA.I. solution seamlessly into existing business workflows while ensuring adherence to cybersecurity protocols, data privacy policies, and regulatory requirements. A smooth and secure deployment positions financial institutions to fully leverage the potential of novel GenA.I. technologies, enabling them to enhance operational efficiency, achieve business objectives more effectively, and deliver superior services to their end clients.

5.2.3. Post-deployment Phase

After deploying the GenA.I. solution, financial institutions should continuously monitor its output, performance, and compliance. It is also important to

raise awareness among end-users and provide them with the necessary skills through training to maximise the uptake. Promoting the successful adoption case to the general public can also be considered to further enhance the financial institution's reputation (see Figure 32).

Figure 32: Post-deployment Phase



Source: Quinlan & Associates analysis

Ongoing Monitoring

Ongoing monitoring involves the continuous assessment and evaluation of GenA.I. solutions, which is essential for financial institutions to mitigate risks related to misinformation, irrelevant references, and output bias or discrimination. By closely and continuously monitoring these solutions, financial institutions can ensure output accuracy, optimise business and technological performance, and maintain compliant and ethical use, ultimately preserving public trust in the technology's reliability and long-term effectiveness.

- **Output Management:** GenA.I. solutions are widely known for their tendency to generate inaccurate outputs (i.e., hallucinations). Therefore, financial institutions should implement procedures to address the potential for incorrect answers provided by GenA.I. to end-users. This includes establishing preventative safeguards and ensuring accountability for outcomes. In our interviews, we observed several practices, such as limiting the GenA.I. solution to respond only within its knowledge base, requiring sources to be provided with answers, and indicating the level of accuracy against the original source. Moreover, OWASP highlighted that financial institutions may conduct regular

reviews of outputs that are being generated by GenA.I. and employ a self-consistency technique. This technique involves using a single prompt on a model multiple times and comparing the responses for consistency. Through this comparison, financial institutions can assess the reliability of the outputs.

- **Solution Performance Monitoring:** Continuous monitoring of the GenA.I. solution's performance is needed, with oversight from both technical and business stakeholders. A robust feedback loop should be established to ensure that observations and user feedback inform future developments and improvements of the GenA.I. solution. An effective logging mechanism should also be implemented to enable ongoing monitoring, taking into consideration the sensitivity of data involved as well as the hardware resource implications. Our interview findings indicate that most financial institutions maintain comprehensive logs of end-user access, input prompts, and generated outputs of GenA.I. solutions. This approach ensures full traceability and explainability, supporting backtracking for developers, output validation for employees, and performance improvement planning for the management team.

- **Compliance Monitoring:** Ongoing compliance monitoring is critical for identifying issues, preserving system integrity, and ensuring the continued safe and responsible operation of deployed solutions. Closely tracking inputs and outputs allows for the detection and mitigation of emerging biases or unintended discriminatory impacts over time, guaranteeing the responsible and ethical use of GenA.I. to secure public trust and maintain the solution's long-term viability.

Training on GenA.I.

To fully activate GenA.I. solutions, financial institutions should equip intended end-users with the necessary knowledge and skills to fully capitalise on these technologies. Training will not only familiarise users with the solution but also increase their confidence, enhancing the overall effectiveness of the solution and leading to better outcomes for the institutions.

- **Training on the General Understanding of GenA.I.:** Financial institutions should educate users on what GenA.I. is, how it works at a high level, and what its benefits and limitations are. This foundational understanding will help users grasp the overall concept of the solution and recognise the risks associated with its use.
- **Training on Specific GenA.I. Solution Usage Methods:** Each GenA.I. solution should come with tailored training for end-users, particularly for more complex and customised internal solutions. Training formats should vary based on the audience. For example, customer-facing GenA.I. solutions may include short instructional videos integrated into the UI, enabling rapid understanding. For internal solutions, financial institutions can look to provide in-person training sessions, online video tutorials, and user manuals.

Marketing to the Public

Lastly, financial institutions should consider sharing their progress in GenA.I. adoption to enhance their public perception and market proposition. Emphasising the commitment to continuously adhere to regulatory principles through a robust governance structure may effectively communicate the organisation's dedication to corporate social responsibility and the public's interests.

Financial institutions can announce their use of GenA.I. through various channels, such as in-app notifications, website banners, and press releases on news websites. This public communication can showcase the organisation's commitment to innovation and its proactive approach to integrating the latest technology for the benefit of both employees and clients.

Conclusion


The responsibilities associated with technology development extend beyond successful deployment, requiring continuous effort and commitment from the organisation. Regular monitoring of performance and compliance is key to ensuring the responsible and ethical use of GenA.I. solutions. Training should focus on enhancing end-user confidence in using these solutions, creating a supportive environment for ongoing adoption. Transparent public communication can further build trust and showcase a financial institution's dedication to innovative yet accountable technology deployment for the benefit of all stakeholders.

5.3. Feedback Loop

One of the key recurring considerations across the deployment value chain is that collecting feedback from the end-user is very important for successfully deployment to ensure smooth operation and continuous improvement but also drives wider adoption within the organisation. The feedback loop consists of three key steps: (1) feedback collection, (2) feedback analysis, and (3) feedback implementation (see Figure 33).

- **Feedback Collection:** During the pre-deployment phase, financial institutions should establish a channel for intended end-users to provide initial feedback, ensuring the solution's accessibility and intuitiveness. Once the solution reaches the PoC stage, it is essential to include a feature at the UI level that allows end-users to give feedback on both the overall performance of the solution and the quality of its outputs. The financial institutions we interviewed understand the importance of collecting feedback and have implemented features that enable end-users to rate the GenA.I. solution on metrics such as response speed, relevance, and accuracy. On the back-end financial institutions must track the

Figure 33: Feedback Loop



Source: Quinlan & Associates analysis

logs generated by the GenA.I. solution, allowing the monitoring team to report significant errors, such as biased responses, so they can be promptly addressed to prevent larger negative impacts.

- **Feedback Analysis:** Collected feedback should be regularly aggregated, cleansed, and analysed to assess both the strengths and weaknesses of the GenA.I. solution. This process helps identify patterns and underlying causes, which are essential for developing future improvement plans.
- **Feedback Implementation:** Based on the analysis collected, GenA.I. solutions should be regularly updated according to improvement plans that build on top of user feedback, with ongoing collaboration between engineers, product managers, and end-users to ensure that new versions effectively address the identified issues. A robust feedback implementation mechanism should be in place to ensure effective and ongoing incorporation of feedback into the GenA.I. solution, such as through back-

end prompt engineering, output filtering, or model augmentation / adaptation. In accordance with ISO / IEC 42001:2023, having these activities or mechanisms in place will enable organisations to improve the effectiveness and performance of their system – an essential element for certification.

5.4. Conclusion

Guided by a robust governance structure, the adoption of GenA.I. solutions in financial institutions requires thorough consideration at both the organisational and solution levels across all stages of deployment.

A thorough impact assessment of these considerations, coupled with proactive planning, can significantly smoothen the deployment process. By doing so, financial institutions can not only mitigate potential risks but also unlock GenA.I.'s full potential, driving innovation, achieving desired business outcomes, and ultimately rendering better end-user experiences.

6. Facilitation Initiatives in Hong Kong

As GenA.I. continues to gain momentum, the financial services industry is actively exploring GenA.I. adoption opportunities, though several roadblocks remain. In addition to the complexities of compliance and the knowledge gaps in practical deployment, financial institutions expressed their concerns about the substantial investment costs,

time, and talent required to successfully deploy GenA.I. solutions.

Recognising these challenges and needs, a number of facilitation initiatives have been introduced in Hong Kong by public entities and public-private partnerships, supporting (1) product launch and (2) product promotion (see Figure 34).

Figure 34: Facilitation Initiatives

	Institution(s)	Facilitation Initiative
 PRODUCT LAUNCH	- The Government of the Hong Kong Special Administrative Region of the People's Republic of China - Cyberport	Three-year A.I. Subsidy Scheme for Computing Power (2024) The Government allocated HKD 3 billion to Cyberport for a three-year A.I. Subsidy Scheme aimed at assisting local universities, research institutes, and businesses in utilising the computing power of the A.I. Supercomputing Centre.
	- Alibaba Hong Kong Entrepreneurs Fund - SenseTime - Hong Kong Science and Technology Park	HKAI Lab Accelerator Programme (2018 - Present) Funded by Alibaba Hong Kong Entrepreneurs Fund and SenseTime under the support of Hong Kong Science and Technology Park, the 12-month programme is dedicated to supporting the commercialisation of A.I. technologies developed in the lab.
	- Hong Kong University of Science and Technology - Alibaba Cloud	MoU on Talent Development and Research Collaboration (2022) Hong Kong University of Science and Technology and Alibaba Cloud will further their collaboration in nurturing local technology talent, collaborating on cutting-edge technology research, and facilitating the research works in the field of A.I., data analytics, and cloud.
	- Cyberport - Cisco	MoU on Establishing an A.I. Lab (2024) Under the Memorandum of Understanding, Cyberport and Cisco agree to jointly establish an A.I. lab at Cyberport, focusing on enhancing A.I. network technologies to improve performance, adaptability, and efficiency to drive innovation in Hong Kong.
	- Hong Kong Investment Corporation - SmartMore	Planned Set-up of A.I. Research Institution and Talent Nurturing (2024) Backed by funding from Hong Kong Investment Corporation Limited, SmartMore is dedicated to establishing Hong Kong's first A.I. Research Institute and partnering with local universities to nurture Hong Kong's young talents in the field of A.I.
	- Hong Kong University of Science and Technology - Alibaba Group	MoU on Establishing Joint Lab for Big Data and A.I. (2024) Hong Kong University of Science and Technology signed a Memorandum of Understanding with Alibaba Group to establish a joint lab on Big Data and A.I. in the next 3 years to foster forward-looking research in innovative applications, enabled by research resources sharing.
	- InvestHK - Financial Services and the Treasury Bureau - Hong Kong Monetary Authority - The Securities and Futures Commission - The Insurance Authority	Hong Kong Fintech Week (2016 – Present) Hong Kong Fintech Week is one of Asia's largest and most influential annual Fintech events. In 2024, it will have an A.I. and Advanced Tech Forum where experts will discuss how A.I. can be leveraged in the financial services sector.
	- Innovation, Technology, and Industry Bureau - Hong Kong Trade Development Council - InvestHK	InnoEX (2023 – Present) InnoEX is a large-scale annual event designed to showcase innovative solutions that empower businesses to capitalise on emerging global technology trends. The previous event featured a diverse array of GenA.I. projects and developments.
 PRODUCT PROMOTION	- Hong Kong Monetary Authority - Cyberport	FiNETech - Into the A.I.verse (2024) The second FiNETech event showcased diverse, innovative A.I. technologies and explored the potential applications of the technologies in areas such as risk management, anti-fraud, customer service, and process transformation.

Source: Alibaba Group, Cisco, Cyberport, Financial Services and the Treasury Bureau, Hong Kong Investment Corporation (HKIC), HKMA, Hong Kong Science and Technology Park (HKSTP), Hong Kong Trade Development Council, Hong Kong University of Science and Technology (HKUST), Innovation, Technology, and Industry Bureau, IA, InvestHK, The Securities and Futures Commission (SFC), SenseTime, SmartMore, The Government of the Hong Kong Special Administrative Region of the People's Republic of China (HKSAR Government), Quinlan & Associates analysis

6.1. Product Launch

In the initial phase, financial institutions explore emerging A.I. models, architectures, and training techniques to evaluate their potential for upcoming projects. Given the rapid evolution of A.I. technologies, comprehensive R&D is required to ensure that GenA.I. solutions are innovative, functional, scalable, and aligned with industry-wide needs. However, the complexity and resource demands of this phase underscore the need for robust support schemes to facilitate successful product launches. Hong Kong's public and private sectors have responded by providing or committing to numerous initiatives to support this process.

To promote the development of Hong Kong's A.I. ecosystem, the HKSAR Government has allocated HKD 3 billion in the 2024-25 budget for a three-year A.I. Subsidy Scheme. This initiative supports local universities, R&D centres, government departments, and A.I.-related enterprises in leveraging the A.I. Supercomputing Centre (AISC) at Cyberport, which is set to reach 3,000 petaFLOPS (equivalent to processing 10 billion images per hour) in capacity by early 2026. Part of the scheme will also be directed towards enhancing the security of the AISC and attracting global A.I. talent and enterprises to the city. To oversee this initiative, the HKSAR Government appointed the A.I. Funding Scheme Committee in August 2024.

Public-private partnerships have also been established to create a favourable product launch environment, enhancing R&D infrastructure and expanding the A.I. talent pool. For example:

- In October 2018, the HKAI Lab, a non-profit initiative funded by the Alibaba Hong Kong Entrepreneurs Fund and SenseTime, and supported by HKSTP, was established to help A.I. companies scale their operations. The lab offers a flagship 12-month Accelerator Programme, allowing participants to access high-performance computing resources, Alibaba's machine-learning platform, and technical support from leading A.I. solution providers. Additionally, the programme provides financial support, advisory services, and valuable networking opportunities.
- In August 2022, Alibaba Cloud and HKUST partnered to launch a joint talent development programme focused on data analytics, cloud computing, and A.I., which includes hands-on workshops, seminars, and internship opportunities, providing students with practical experience at Alibaba Cloud.
- In May 2024, Cyberport and Cisco signed the Memorandum of Understanding (MoU) that aims to enhance A.I. networks and cybersecurity in Hong Kong. As part of this collaboration, the two entities will create an A.I. lab at Cyberport to improve A.I. network technologies and drive innovation.
- In June 2024, the HKIC invested in SmartMore, a home-grown A.I. firm, to accelerate technological advancements. SmartMore plans to establish an A.I. Research Institute and collaborate with local universities to offer Master's and Doctoral programmes in A.I., building a pipeline of highly skilled professionals to drive future innovations.
- In July 2024, HKUST and the Alibaba Group signed an MoU that outlines their plans to create a Joint Lab on Big Data and Artificial Intelligence in the next three years, focusing on researching emerging technologies such as GenA.I.

The collaborative efforts of the HKSAR Government, quasi-government agencies, and private sector organisations are instrumental in advancing A.I. in Hong Kong. Through targeted financial support, comprehensive talent development programmes, and strategic joint research ventures, these measures ensure that researchers and enterprises have access to essential resources needed to develop GenA.I. solutions. Importantly, many of these support schemes alleviate concerns raised by institutions about upfront costs and resource needs, facilitating long-term innovation.

6.2. Product Promotion

Effective deployment and promotion are crucial to the success of GenA.I. solutions post product launch.

Hong Kong FinTech Week is one of Asia's largest and most influential flagship Fintech events, spotlighting cutting-edge developments, including A.I. In 2023, the event featured over 300 speakers, 540 sponsors and exhibitors, and attracted more than 30,000 attendees, with 5 million views from over 90 economies. The event includes exhibitions, keynote speeches, panel discussions, workshops, networking opportunities, and business matchmaking sessions. The 2024 Hong Kong FinTech Week is set to take place in late October, with a forum dedicated to A.I. and Advanced Tech.

Another large-scale event is the InnoEx, which showcases innovative solutions aligned with emerging global technology trends. The 2024 event featured 458 exhibitors from 13 countries and attracted over 39,000 buyers, with A.I. & Robotics among the top trending topics. Like Hong Kong FinTech week, the event provides an invaluable platform for institutions to promote their A.I. solutions and engage with a global audience.

Most recently, the HKMA's FiNETech2 event gathered over 300 professionals from the banking, securities, and insurance sectors, as well as representatives from the government and industry associations, to explore the applications of A.I. and GenA.I. in risk management, anti-fraud, customer service, and process transformation. 14 A.I. providers showcased diverse, innovative solutions through booth exhibitions and demonstration sessions. Experts from academia and industry also shed light on GenA.I.'s impact on A.I.-driven digital experiences, highlighting associated benefits and risks.

Large-scale events not only provide a venue for showcasing innovations but also play a crucial role in the GenA.I. commercialisation process by fostering valuable connections and dialogue with industry leaders, potential partners, and end-users. Participation in these events also allows institutions to stay informed about the latest developments and trends in the evolving GenA.I. landscape, ensuring their solutions remain aligned with evolving market needs, crucial for driving sustained commercial success.

7. Way Forward

As GenA.I. continues to evolve, it is crucial for both technology providers and users to stay informed about the latest risk landscape as well as policy support initiatives. Keeping abreast of these developments enables stakeholders to accelerate the development and deployment of A.I. solutions, particularly in the financial sector. Efficiently leveraging government resources and platforms can unlock new opportunities for research and business partnership, especially for financial institutions looking to tailor customer-facing A.I. solutions to local market preferences. Furthermore, the public sector can draw on the policy support being rolled out in other jurisdictions to enhance local efforts that foster ongoing innovation in GenA.I., including providing opportunities for financial institutions and solution providers to scale their innovations.

Recognising the importance of A.I. adoption among the market participants, the HKMA has consistently shared its vision through a series of whitepapers, including *Reshaping Banking with Artificial Intelligence* (2019), *Artificial Intelligence in Banking* (2020), and *Artificial Intelligence and Big Data in the Financial Industry* (2021). These publications highlight the transformative potential of A.I. within the financial services sector and address key themes that are highly relevant to the evolution of GenA.I. solutions. The HKMA also acknowledges the new risks introduced by this emerging chapter in A.I. and thus underscores the need for a more interactive and iterative supervisory approach to assisting financial institutions in Hong Kong in navigating the adoption of the technology, particularly in the areas of risk management, anti-fraud, and customer interface.

In response to these emerging challenges and opportunities, the GenA.I. Sandbox comes at the right time to accelerate the adoption of advanced A.I. technologies within Hong Kong's Banking sector. The controlled and risk-managed environment allows banks to test the viability of their business models, collect valuable data, and gather user feedback. Additionally, participants can directly engage with the HKMA to gain early supervisory feedback, ensuring alignment with future regulatory guidance.

As we look to the future, GenA.I. is set to play a transformative role in shaping the evolution of the global financial services industry, driving both innovation and operational efficiency. However, the journey towards widespread adoption will demand ongoing collaboration between regulators, financial institutions, and technology providers. This partnership will be essential for ensuring that the integration of A.I. technologies is both forward-looking and responsibly managed.

The HKMA remains committed to identifying and implementing policy-level support measures to promote the responsible adoption of GenA.I. across the banking industry. By fostering an environment of compliance, innovation, and safeguards, the HKMA aims to ensure that banks in Hong Kong can leverage the transformative potential of GenA.I. while maintaining the financial stability, security, and trust that are the cornerstones of the industry.

In collaboration with other financial regulators including the SFC, the IA and the MPFA, the HKMA will continue to develop and refine policies that promote responsible technological innovation. Initiatives, such as the GenA.I. Sandbox, are only the beginning. These efforts are designed to foster a secure and controlled environment where financial institutions in Hong Kong can confidently explore and adopt advanced A.I. technologies.

Through proactive collaboration and open communication between the public and private sectors, the HKMA will continue to support Hong Kong's position as one of the world's leading international financial hubs. We believe that by working together, we can achieve a future where financial institutions not only benefit from the advantages of GenA.I. but also contribute to a responsible, sustainable, and innovative financial ecosystem.

8. Appendix

8.1. List of Abbreviations

- **A.I.:** Artificial Intelligence
- **AIDA:** Artificial Intelligence and Data Analytics
- **AISC:** A.I. Supercomputing Centre
- **API:** Application Programming Interface
- **CPU:** Central Processing Unit
- **ESMA:** European Securities and Markets Authority
- **EU:** European Union
- **FCA:** Financial Conduct Authority
- **GenA.I.:** Generative Artificial Intelligence
- **GPU:** Graphics Processing Unit
- **HKIC:** Hong Kong Investment Corporation
- **HKMA:** The Hong Kong Monetary Authority
- **HKSAR Government:** The Government of the Hong Kong Special Administrative Region of the People's Republic of China
- **HKSTP:** Hong Kong Science and Technology Park
- **HKUST:** Hong Kong University of Science and Technology
- **IA:** The Insurance Authority
- **IDS:** Intrusion Detection System
- **IMDA:** Infocomm Media Development Authority
- **IT:** Information Technology
- **KYC:** Know-Your-Customer
- **LLM:** Large Language Model
- **LoRA:** Low-Rank Adaptation
- **MoU:** Memorandum of Understanding
- **MPFA:** The Mandatory Provident Fund Schemes Authority
- **NIST:** National Institute of Standards and Technology
- **OWASP:** Open Web Application Security Project

- **PII:** Personally Identifiable Information
- **PoC:** Proof-of-Concept
- **RAG:** Retrieval-Augmented Generation
- **RM:** Relationship Manager
- **SAR:** Special Administrative Region
- **SEO:** Search Engine Optimisation
- **SFC:** The Securities and Futures Commission
- **TPU:** Tensor Processing Unit
- **U.K.:** The United Kingdom
- **UI:** User Interface
- **US:** United States
- **UX:** User Experience

8.2. Glossary of Key Terms

- **Artificial Intelligence:** A family of technologies that involve the use of computer programs and machines to mimic the problem-solving and decision-making capabilities of human beings.
- **Central Processing Units:** Computational engines in computers that execute instructions sequentially and perform complex calculations to manage and coordinate computer systems.
- **Corpus data:** A vast dataset used by A.I. solution providers to begin training the model.
- **Decoder-only models:** Models that generate text by predicting the next word in a sequence based on the context of previous words.
- **Deep learning model:** Models that use modern neural networks, characterised by having many layers, enabling them to learn complex patterns from data.
- **Diffusion models:** A type of generative model that produces images, audio, and video outputs.

- **Encoder-decoder models:** Models that use encoders to process and transform input text into a representation and decoders to generate the corresponding output text.
- **Encoder-only models:** Models that encode the input text into a fixed-length representation, capturing contextual information for further processing or analysis.
- **Feed-forward mechanism:** The process where the input data passes through multiple layers of a neural network, with each layer performing calculations and applying distinct weights to the inputs.
- **Generative Artificial Intelligence:** A category of A.I. that excels at generating novel and innovative outputs across various domains by utilising generative models.
- **Graphic Processing Units:** Specialised hardware accelerators that utilise numerous smaller, highly efficient processing cores to process data in parallel, making them particularly effective for tasks such as graphics rendering and complex computations.
- **Large Language Model:** A type of generative models that can produce contextually relevant language and code outputs, leveraging vast amounts of training data to understand and generate text.
- **Model architectures:** Design configurations and structural frameworks of A.I. models that dictate how they process information and generate outputs, influencing their performance and capabilities.
- **Nodes:** Individual computation units within an A.I. model, responsible for processing data and contributing to the overall functionality of the network.
- **Parameters:** Metrics that determine how input signals are transformed as they propagate through the network, including weights that regulate the strength of connections and biases to adjust the output independently of the input.
- **Processing core:** Units that retrieve instructions from memory in the form of digital signals, decode them, and execute computations and logical functions using logical gates.
- **Prompt:** A query or request for specific output, serving as an instruction to an application interface powered by an LLM.
- **Proof-of-Concept:** A demonstration or prototype that is created to validate an idea, concept, or theory, helping organisations make informed decisions about moving forward with a project.
- **Reinforcement data:** Data used for reinforcement learning.
- **Reinforcement training:** A training method in which models interact with their environment, receiving rewards for desirable actions and penalties for undesirable ones, allowing for iterated refinement and improvement in task performance over time.
- **Retrieval-augmented generation:** An approach that enables an A.I. model to access data beyond its original training corpus, integrating external, real-time information to generate more accurate, relevant, and up-to-date responses.
- **Self-attention:** The mechanism that enables a model to evaluate the importance of each word in a sentence relative to every other word.
- **Semantic meaning / understanding:** The meaning of words, phrases, and sentences in a language.
- **Sparse models:** Models where only a small fraction of parameters is non-zero, meaning that most of the parameters are effectively unused or set to zero to improve computational efficiency and interpretability.
- **Supervised training:** A training method in which the model receives labelled training data where the correct output is already known so the model can learn to map inputs to their corresponding outputs.
- **Tensor Processing Units:** A.I.-specific accelerators that are designed to optimise complex computations involved in training and running A.I. models.
- **Transformers:** The standard architecture for building LLMs that leverages mechanisms like self-attention to process data in parallel, allowing for efficient training and effective handling of sequential information.
- **Unsupervised training:** A training method where the model autonomously explores a dataset to identify underlying patterns and structures without the use of predefined outputs or labels.

8.3. List of Reference

- **Cyberspace Administration of China, the National Development and Reform Commission, the Ministry of Education, the Ministry of Science and Technology, the Ministry of Industry and Information Technology, the Ministry of Public Security, and the National Radio and Television Administration:** Interim Measures for the Administration of Generative Artificial Intelligence Services (2023)
- **Digital Policy Office of the Hong Kong Government:** Ethical Artificial Intelligence Framework (2024)
- **European Data Protection Board ChatGPT Taskforce:** Report of the Work Undertaken by the ChatGPT Taskforce (2019)
- **European Securities and Markets Authority:** Public Statement on the Use of Artificial Intelligence in the Provision of Retail Investment Services (2024)
- **European Parliament and the Council of the European Union:** Artificial Intelligence Act (2024)
- **European Union Agency for Cybersecurity:** Multilayer Framework for Good Cybersecurity Practices for AI (2023)
- **Financial Conduct Authority:** AI Update (2024)
- **High-Level Expert Group on Artificial Intelligence:** Ethics guidelines for trustworthy AI (2019)
- **Hong Kong Academy of Finance & Hong Kong Institute for Monetary and Financial Research:** Artificial Intelligence in Banking (2020), Artificial Intelligence and Big Data in the Financial Industry (2021)
- **Hong Kong Monetary Authority:** Consumer Protection in respect of Use of Generative Artificial Intelligence (2024), High-level Principles on Artificial Intelligence (2019), Consumer Protection in respect of Use of Big Data Analytics and Artificial Intelligence by Authorized Institutions (2019), Reshaping Banking with Artificial Intelligence (2019)
- **Infocomm Media Development Authority and Personal Data Protection Commission:** Model Artificial Intelligence Governance Framework Second Edition (2020)
- **Infocomm Media Development Authority:** Model AI Governance Framework for Generative AI: Fostering a Trusted Ecosystem (2024)
- **Information Commissioner's Office and The Alan Turing Institute:** Explaining Decisions Made with AI (2022)
- **Information Commissioner's Office:** Generative AI: Eight Questions that Developers and Users Need to Ask (2023)
- **Insurance Authority:** Conduct in Focus - Chatting about Chatbots and AI (2023)
- **Intellectual Property Department:** Copyright and AI, Public Consultation Paper (2024)
- **Monetary Authority of Singapore:** Emerging Risks and Opportunities of Generative AI for Banks A Singapore Perspective (2023)
- **National Cyber Security Centre:** Guidelines for Secure AI System Development (2023)
- **National Technical Committee 260 on Cybersecurity:** Basic Security Requirements for Generative Artificial Intelligence Service (2024)
- **National Institute of Standards and Technology:** Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (2024)
- **National New Generation Artificial Intelligence Governance Specialist Committee:** Ethical Norms for New Generation Artificial Intelligence (2021)
- **Office for Artificial Intelligence and Department for Science, Innovation & Technology:** A Pro-innovation Approach to AI Regulation (2023)
- **Office of the Government Chief Information Officer:** Ethical Artificial Intelligence Framework (2023)
- **Office of the Privacy Commissioner for Personal Data:** Artificial Intelligence: Model Personal Data Protection Framework (2024), Guidance on the Ethical Development and Use of Artificial Intelligence (2021)
- **People's Bank of China:** Guidance on Information Disclosure for Financial Applications Based on Artificial Intelligence Algorithms (2023)
- **Personal Data Protection Commission:** Advisory Guidelines on Use of Personal Data in AI Recommendation and Decision Systems (2024)
- **The White House:** Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (2023)

8.4. Acknowledgements

This paper has greatly benefitted from the contributions of various external collaborators. We thank the Insurance Authority, the Mandatory Provident Fund Schemes Authority, the Hong Kong Institute for Monetary and Financial Research, various divisions and departments of the Hong Kong Monetary Authority, including the Communications Division, the Banking Supervision Department, the Banking Conduct Department and the Fintech Facilitation Office for their comments and suggestions.

We would also like to thank Quinlan & Associates and KPMG China for coordinating and facilitating interviews with various market participants, including technology providers, banks, securities, and insurance firms, and for their support in the research and drafting of this paper.

We wish to extend our gratitude to Amazon Web Services Hong Kong Limited, Bank of China (Hong Kong) Limited, Bank of Communications (Hong Kong) Limited, Citibank (Hong Kong) Limited, FIL Investment Management (Hong Kong) Limited, Hang Seng Bank, Limited, Hongkong and Shanghai Banking Corporation Limited (The), HSBC Life (Asia) Limited, J.P. Morgan Securities (Asia Pacific) Limited, JPMorgan Chase Bank, National Association, Microsoft Hong Kong Limited, Moody's Asia Pacific Ltd, OneDegree Hong Kong Limited, PAO Bank Limited, Standard Chartered Bank (Hong Kong) Limited, and Zurich Insurance (Hong Kong) for their inputs for this paper.