

GenAI as a catalyst for lasting change in financial compliance operations

A collaboration between Deloitte and Hummingbird

Introduction: compliance & legacy systems

There's no shortage of technological innovation in the financial sector. But in practice areas that prioritize security, consistency, and continuity, the implementation of new technology is extraordinarily difficult. Change management concerns and the highly embedded nature of legacy systems combine to form a series of roadblocks,

blocking the deployment of newer, more cutting-edge technology into financial institutions' core systems.

Financial compliance, in particular, often struggles to adapt to new technologies. As a practice area, compliance is defined by its need for effective risk mitigation, audit-readiness, and predictable operating procedures. Consequently, compliance leaders are necessarily

more concerned with the safety and consistency of their operations than they are with pushing the technological envelope. ➤

A paradigm shift: Generative AI in banking operations

Sometimes, however, a technology comes along that has the potential to transform even these formerly impenetrable practice areas. It is our belief that large language models (LLMs) and other types of generative AI (GenAI) are a paradigm-shifting technology, with the potential to not only transform work involving sensitive banking operations but also the barriers presented by legacy systems. Many of the technologies in financial institutions are mission-critical, and there is a general reluctance to alter them due to the fear of disrupting essential functions. As a tool, however, GenAI stands out due to its flexibility; in many instances, it can be layered on top of existing systems, creating new business capabilities without significantly altering mission-critical operations.

As with any new technology, fears about potential inaccuracies are commonplace amongst corporate IT and especially compliance teams. With GenAI, these concerns (on model hallucination, lack of transparency, and data protection) remain at the forefront of any discussion about the technology's potential implementation. Unfortunately, these concerns often lead financial compliance teams to the conclusion that GenAI tooling is incompatible with their internal systems, when in fact it is possible to address and mitigate these concerns through strategic and compliant deployment.



What makes the use of GenAI in financial institutions different?

GenAI tools feature several fundamental advantages over previous types of “transformative” technologies in the financial space. First and foremost, GenAI solutions are exponentially faster to build and implement than their predecessors, dramatically shortening the go-to-market timeline. And unlike previous banking operations technologies – like Robotic Process Automation (RPA), for example – they do not require extensive involvement or continued maintenance from engineers and data scientists. This reduces the need for a highly technical background among users, democratizing the use of these tools within financial institutions.

Within compliance, these advantages are helpful, if not a sure-fire ticket to overcoming the barriers to implementation. Operational challenges are only half the battle, as compliance departments must also reckon with intense regulatory scrutiny, which can be highly critical of any bank’s attempt to do anything new or different regarding its risk mitigation practices. But as previously noted, we believe that GenAI tools have the flexibility and adaptability required to overcome these challenges, so long as the financial institution follows a structured, compliance-first approach.



Challenges of GenAI in compliance

The key challenges for compliance teams looking to incorporate AI-based tools into their investigative and operational workflows are:

01

Data privacy concerns

Creating an effective AI tool often requires training and calibrating it on large amounts of relevant data. Given the amount of sensitive financial information and personally identifiable information (PII) involved in compliance work, concerns about privacy and data security are well-founded.



02

Interoperability issues

Making sure that new AI tools are compatible with existing systems and processes may pose challenges, especially when they are being built upon older, legacy infrastructure.



03

Explainability and transparency

Ensuring that AI decisions are explainable and transparent minimizing the risk of hallucinations is essential for compliance audits and regulatory scrutiny – something many “out of the box” GenAI solutions can’t do.



04

Accuracy and auditability

Regulatory requirements require that compliance work be accurate and maintain a clear record of decisions and actions taken, meaning that any AI technology used must be able to produce results that are both accurate and accountable.



05

Vendor selection

Choosing the right vendor for AI compliance solutions requires finding a supplier who combines deep knowledge of AI technologies with a proven track record for helping companies navigate the complexities of the regulatory landscape.



Scalable AI solutions

In the context of financial compliance, it is essential to build efficiently scalable AI solutions. Compliance teams increasingly rely on AI to process large amounts of data, comply with regulatory requirements, and stay operationally efficient. Scalable AI solutions allow the system to adjust to growing demands when the workload increases.

This can be due to more intensive data analysis or a higher number of parallel user requests. The system can then scale either horizontally by increasing the number of server nodes to handle more parallel user requests or vertically by allocating more processing units to the existing servers to compute complex calculations more efficiently. The ability to handle many complex user requests smoothly is of essential importance.

Our approach uses a flexible modular design, which allows different components to work independently and together as needed. This design enables us to adjust our applications to fast-paced environments such as the regulatory one. Furthermore, our flexible approach ensures that updates and new functionalities can be rolled out more easily without affecting the overall system, thereby guaranteeing consistency and frictionless compliance operations.

The scaling of AI solutions does not solely rely on processing more data or users; it is about keeping a continuously developing system compliant. A scalable system means that financial institutions can deploy updates and new features without the need for costly system adjustments or disruptions to existing compliance processes. This approach

allows institutions to proactively adjust to changing regulatory environments while staying operationally efficient and resilient.

With the increasing number of AI systems in the financial compliance context, it is essential to build scalable systems while relying on highly flexible and modular solutions. This approach allows financial institutions to comply with today's requirements and prepare for tomorrow's challenges.

Strategies for effective (and compliant) GenAI implementation in the financial industry

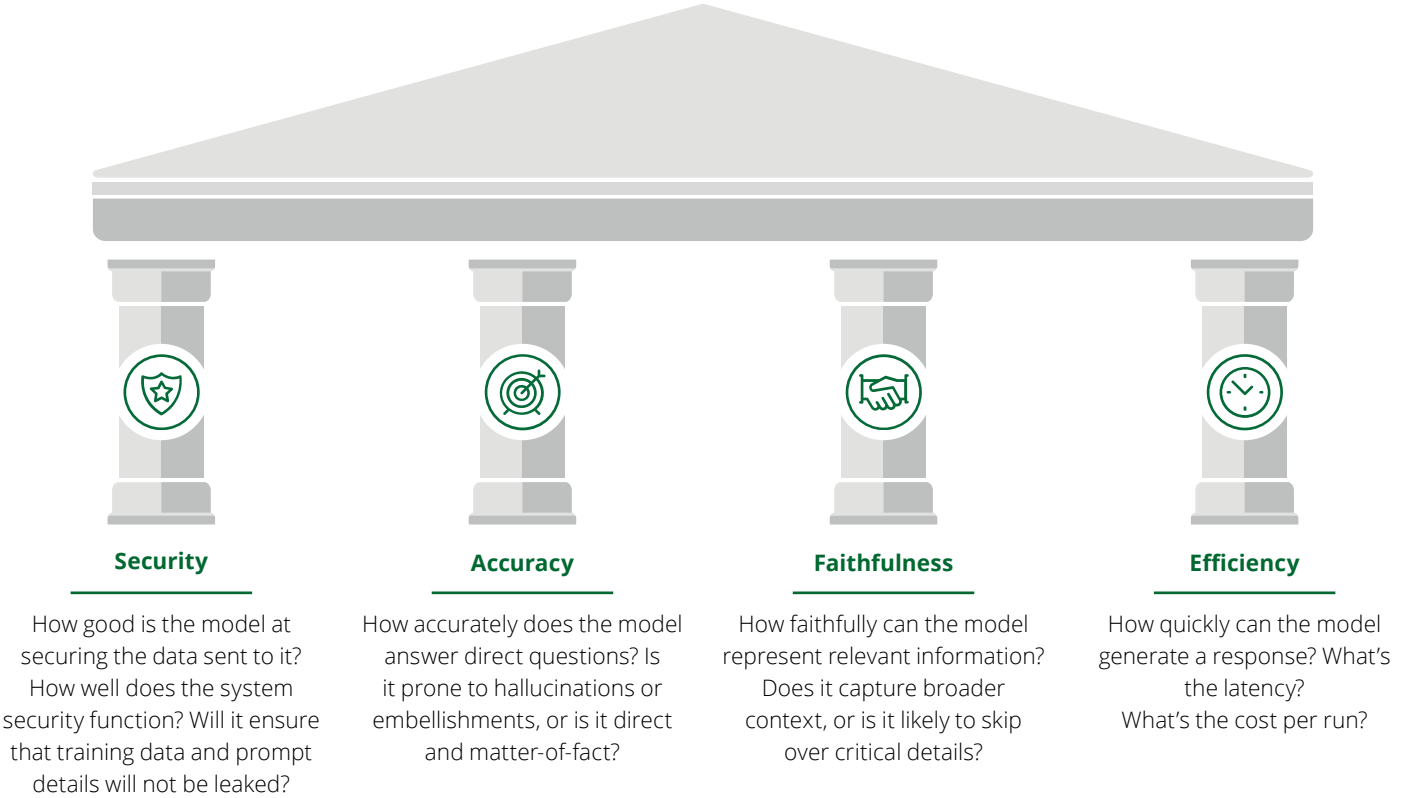
While we have so far focused on the need for scalability and modularity in a compliance-oriented genAI approach, it's important to factor implementation into the equation. Regarding implementation and change management, financial institutions should follow a strategic, structured approach. We recommend a strategy that invests time and effort before, during and after implementation.

Before implementation
Create or adopt frameworks that guide GenAI requirements and success metrics

These frameworks will help guide you through the vendor selection process as well as help define the use cases for GenAI feature applicability.

Example: The Hummingbird S.A.F.E. framework for AI-model selection

Hummingbird has developed a framework to help both regtech firms and financial institutions decide which LLM model to use when developing their AI capabilities for compliance work, naming it the "S.A.F.E. framework" after its four pillars:



During implementation

Keep AI applications focused and build iteratively

By limiting the number of GenAI use cases that are being developed rather than trying to apply it directly to your entire operation, you can configure, iterate and refine your solution – making it more effective in both the short and long term.

Example: Hummingbird targeted AI feature development

Hummingbird's philosophy to AI feature development is "build features that solve compliance problems first, and employ AI second." To this end, they have been rigorous in their focus on developing AI-powered tools that answer a specific need. Their first feature, AI File Summarization, provided compliance-oriented summaries of case documents at the push of a button, giving CDD teams a huge boost in productivity.

Next, Hummingbird-focused on AI Suspicious Activity Report (SAR) narrative generation. Early experiments with this use case taught the Hummingbird team several valuable lessons, and informed their strategy. Rather than aim to use AI to create a "foolproof" narrative, Hummingbird's Narrative Generation tool (now in beta) focuses on delivering a high-quality first draft – informed by case data – for editing and review by the human investigator.

In tandem with its narrative generation feature, Hummingbird created a complimentary tool: AI narrative validation. This feature turns the power of AI not on creating the content of a narrative, but on reviewing the existing narrative for accuracy and completeness. It is able to highlight any discrepancies between the content of the narrative and the case data, as well as to suggest additional potential items for inclusion.

Taken together, Hummingbird's growing selection of AI tools provides evidence of their compliance-first approach to AI feature development. This is proven out in the simple fact that the strategic focus of each tool is always clearly grounded in the same goal: allowing compliance teams to drive efficiency while also maintaining the highest standards of accuracy and attention-to-detail.

Example: Deloitte PolicyBot development

In its constant endeavor to design a GenAI application tailored for addressing compliance requirements, Deloitte developed a unique LLM expert system. Although creating an application capable of addressing complex compliance issues was a challenging process, the development team's consistent focus on domain-specificity yielded an effective solution. One example that became clear throughout the development of the tool due to iterative testing is the importance of finetuning LLM models. The meaning of "monitoring" and "screening" in another context is rather similar whereas for the anti-financial crime domain the terms describe entirely different topics. Using said iterative approach, the team fashioned a tool that reads, learns, and applies domain-specific content. Now, with curated, trustworthy sources and thorough tuning, the application delivers insightful compliance knowledge. While the main goal has always been to assist compliance departments in addressing complex queries, the process has produced an effective LLM tool that aligns and learns specifics, a benefit for any high precision work environment. By designing this tool, the Deloitte team not only continues to enhance its LLM system knowledge underlining the importance of curation of sources, traceability through quotations, source highlighting options, and a robust audit trail and access management system

but the team also developed capabilities for unique support to teams needing accuracy at a swift pace.

Both examples demonstrate how GenAI can be of enormous value for compliance use cases, but before blindly deploying solutions for the sake of using new technologies, decision makers should thoroughly think about what should be achieved, which potential pitfalls should be addressed and IT experts should collaborate closely with domain experts in order to achieve the expected results.

After implementation

Continuous monitoring and assessment; cross-system integration

No technology implementation should be considered "finished" just because it is live and in production. Having a process in place for regularly reviewing performance analytics for LLM-backed tools is essential.

Hummingbird and Deloitte: AI tools designed for compliance

Hummingbird and Deloitte possess a mutual understanding of GenAI requirements and display a conscientious approach toward their integration strategies. Hummingbird, a purveyor of cutting-edge compliance operations and investigative tools, has been a game-changer for financial institutions. Thanks to adjustable process automation, customizable workflows, and a suite of compliance-grade AI tools, it enables these institutions to accomplish more with fewer resources. Founded by technologists, product designers, and ex-regulators, Hummingbird prioritizes “compliance first” approach to AI feature development. With AI tools adept at document summarization, case prep, and narrative validation, they have rapidly created an essential suite of AI features for compliance teams. In

parallel, Deloitte brings a rich compliance and anti-money laundering expertise to the table. With hands-on experience in GenAI use cases and a finger on the pulse of the market and its vendors, they remain at the forefront. In addition to their Anti-Financial Crime Team, there is a large pool of GenAI experts who strive to keep abreast of the latest technological information, making them a formidable player in the field.

Conclusion

In order to operate successfully, financial institutions – and in particular compliance teams – must hold themselves to exacting standards for accuracy, efficiency, transparency, and interoperability. These demanding operational standards have a downside in that they often prevent new technology (whose implementation is typically costly, disruptive, and time-consuming) from reaching deep into

financial institutions’ tech stack. But GenAI technology, with its wide range of use cases and cross-system adaptability, represents a shift in thinking – a way to bridge the gap from legacy system to cutting-edge technology.

While there are well-founded concerns surrounding the use of AI in sensitive banking operations (specifically related to accuracy and model transparency), these reservations can be successfully overcome through careful planning, targeted application, and by selecting vendor-providers with a demonstrated track record in both AI feature development and the financial industry.



If you want to learn more



Martin Hirtreiter

Deloitte | Partner
Anti-Financial Crime Advisory
mhirtreiter@deloitte.de



Dr. Robert Schmuck

Deloitte | Director
Anti-Financial Crime Advisory
rschmuck@deloitte.de



Maria Antonia Flosbach

Deloitte | Senior Manager
Anti-Financial Crime Advisory
mflosbach@deloitte.de



Matt van Buskirk

Hummingbird | Co-Founder & CEO
Regulatory
matt@hummingbird.co

Deloitte.

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited (DTTL), its global network of member firms, and their related entities (collectively, the “Deloitte organization”). DTTL (also referred to as “Deloitte Global”) and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/de/UeberUns to learn more.

Deloitte provides industry-leading audit and assurance, tax and legal, consulting, financial advisory, and risk advisory services to nearly 90% of the Fortune Global 500® and thousands of private companies. Legal advisory services in Germany are provided by Deloitte Legal. Our people deliver measurable and lasting results that help reinforce public trust in capital markets, enable clients to transform and thrive, and lead the way toward a stronger economy, a more equitable society and a sustainable world. Building on its 175-plus year history, Deloitte spans more than 150 countries and territories. Learn how Deloitte’s approximately 457,000 people worldwide make an impact that matters at www.deloitte.com/de.

This communication contains general information only, and none of Deloitte GmbH Wirtschaftsprüfungsgesellschaft or Deloitte Touche Tohmatsu Limited (DTTL), its global network of member firms or their related entities (collectively, the “Deloitte organization”) is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser.

No representations, warranties or undertakings (express or implied) are given as to the accuracy or completeness of the information in this communication, and none of DTTL, its member firms, related entities, employees or agents shall be liable or responsible for any loss or damage whatsoever arising directly or indirectly in connection with any person relying on this communication. DTTL and each of its member firms, and their related entities, are legally separate and independent entities.