

Financial Services and Generative AI: Navigating a New Era of Innovation

How Financial Services Firms
are Embracing — and Governing
— Generative AI



Table of Contents

- Forward** 3
- Chapter 1: Making Strategic Decisions About Generative AI and Balancing Regulatory Risks** 4
 - How are firms thinking about generative AI today? 5
 - How are firms evaluating the benefits and risks of generative AI? 8
 - How are stakeholder perspectives integrated into generative AI governance and risk management practices? 12
- Chapter 2: Regulatory and Risk Implications: How to Be Ready** 14
 - What methods are being used to identify, assess, and prioritize generative AI risks? 15
 - How are firms assessing the impact of generative AI upon specific regulatory obligations? 17
 - How are firms monitoring developments related to industry standards? 20
- Chapter 3: Generative AI and the Impact on Compliance** 22
 - How are risks being translated into actionable policies for managing generative AI risks? 23
- Chapter 4: Governance, Accountability and Model Safety** 25
 - How are organizations implementing generative AI governance? 26
 - How are firms providing due diligence on existing applications that are now embedding generative AI? 28
 - How are firms evaluating and selecting specific generative AI models? 30

Forward

Generative AI has been unleashed upon financial services, potentially disrupting how firms work, arrive at critical decisions, and interact with the market. It is simultaneously under-hyped and over-hyped, with immense top-down pressure exerted by firm leadership to harness its capabilities, while users are actively experimenting with use cases that are within and outside of existing compliance controls and ambiguous regulatory obligations. Caught in the middle are compliance officers, who are in an uncomfortable position – not to say “No” to the use of generative AI, but to help guide the firm on “How.”

This guide summarizes how compliance and other risk stakeholders can support their firm’s use of this transformative technology within communications and collaborative technology infrastructures while setting the appropriate controls and guardrails to help mitigate the amplification of existing risks and those we are beginning to understand. It features commentary from industry experts, representing a diverse set of business, technological, legal, and human risk disciplines — reflecting the breadth of impact generative AI thrusts upon business and technology objectives.





Chapter 1

Making Strategic Decisions About Generative AI and Balancing Regulatory Risks

How are financial firms thinking about generative AI today?

Financial services firms are approaching generative AI with a mixture of enthusiasm and caution, recognizing its transformative potential while acknowledging the complex regulatory landscape inherent to the industry. The industry has a mixed understanding of the opportunities and challenges this cutting-edge technology presents – which is evolving rapidly.



We're seeing a fairly rapid uptick in interest and adoption for generative AI across numerous enterprise use cases.

Jon Chan, Senior Managing Director, FTI Consulting

Nearly all firms are adopting a phased approach to generative AI adoption. Short-term projects will improve internal efficiency and reduce operating costs, leveraging lower-risk use cases to gain experience and surface obstacles. The intermediate-term goals will shift towards transformative projects such as improving customer service, and longer-term objectives will likely center on scaling the business and uncovering new revenue streams. ([SIFMA & Deloitte Virtual Forum: Generative AI And AI Risk Management](#))

Internally: Organizations seek to improve the efficiency of manually intensive tasks and functions, such as automating the search and retrieval of information, summarizing meetings and documents, strengthening risk management capabilities, and bolstering fraud prevention measures. This internal focus reflects the current industry pressures to increase efficiency, reduce operational risks and improve organizational performance through innovative technologies.

Externally: Firms are beginning to consider client-facing use cases, such as AI-driven customer service solutions, personalized financial advice platforms, and product recommendation systems. These initiatives aim to enhance user experience, improve service delivery, and potentially create new revenue streams. However, each externally facing use case intersects with existing financial services regulatory obligations, which has caused firms to pursue these use cases more cautiously.



It's not the tool you use; it's what these tools could do. The same worries that people have about generative AI were applied to machine learning on structured data around issues like discrimination in consumer lending.

Matthew Bernstein, Information Governance Strategist, MC Bernstein Data

More firms demand that generative AI projects have hard ROI objectives and ties to key initiatives to earn corporate funding and support. As noted in the [Global Insight Report by Citi](#), generative AI “can create the opportunity for innovation and improved quality of life. However, it can also create losers, especially in the short run.”

What is true across all use cases is the rethinking of the human-to-AI collaboration model. Human judgment and quality control remain integral to the process. Defining the specific responsibilities for human co-pilots will be crucial in areas subject to regulatory scrutiny, as well as those use cases that potentially expose intellectual property or raise information security or data privacy risks.

The state of AI regulation also plays a significant role in shaping generative AI strategies, particularly for multinationals, as AI knows no borders. Firms are developing implementation plans that account for both current and anticipated regulations, such as the recently enacted [EU AI Act](#). This proactive stance includes identifying potential “high-risk” AI applications early and establishing robust governance structures and documentation practices.

Rodrigo Madanes from EY noted at the [Reuters AI Momentum](#) event, “Think of generative AI as an example of a consumer application now entering the workforce. You try to skate where the hockey puck is going.” This mentality underscores the innovative mindset that financial services firms are adopting as they assess the possibilities created by generative AI.



Emerging Best Practices

- 1. Balanced dual-focus strategy:** Develop and implement a comprehensive generative AI strategy that inclusively addresses external (client-facing) and internal (operational) use cases, ensuring resources and attention are appropriately allocated between quick wins and transformational, longer-term projects.
- 2. Targeted implementation in high-value areas:** Focus generative AI investment in areas where it can deliver time-to-value, such as customer service, large-scale data analysis, compliance review, and employee support tools.
- 3. Human-AI collaboration model:** Adopt a model that combines generative AI automation with human oversight. This approach expands coverage of critical activities while ensuring that human judgment and quality control remain integral to the process, particularly in areas subject to regulatory scrutiny.

- 4. Regulatory-aware deployment strategy:** Develop deployment strategies that account for regulatory fluidity. This includes earmarking potential “high-risk” AI applications and preparing to meet strict obligations before bringing AI-driven products to market, especially in areas like credit scoring.
- 5. Continually scan employee-driven demand:** New use cases and generative AI tool innovations continue to arrive at a nearly exponential rate. Compliance and risk teams must continuously engage with user groups to assess new opportunities and current controls and minimize duplication of efforts across the firm.



How are firms evaluating the benefits and risks of generative AI?

Financial services firms are quickly evolving from organic to methodical approaches to evaluate the benefits and risks of generative AI for prospective use cases. While many organizations are in the early stages of learning from generative AI exploration and pilot projects, they are acutely aware of the associated risks and regulatory obligations.

A holistic approach to AI governance

Many firms are establishing AI governance councils to guide the evaluation process and risk analysis. These bodies help ensure that generative AI initiatives align with organizational strategies, comply with regulations and adhere to ethical standards.

Firms are also engaging diverse stakeholders to carefully consider potential generative AI uses. This collaborative approach involves internal teams from various departments, including compliance, legal, data science and business



units, as well as external advisers with specialized expertise in generative AI, representing both the human and data science elements.

Many firms are implementing holistic evaluation processes that examine potential generative AI use cases and associated risks across critical business functions, including IT, information governance, privacy, data management, legal, and compliance risk management. These frameworks typically consider multiple factors, such as:

- Business value and impact
- Performance objectives and KPIs
- Security, privacy and IP risk assessment
- Third-party and vendor risk assessment
- Development timeframe and deployment cost and complexity
- Ongoing support requirements

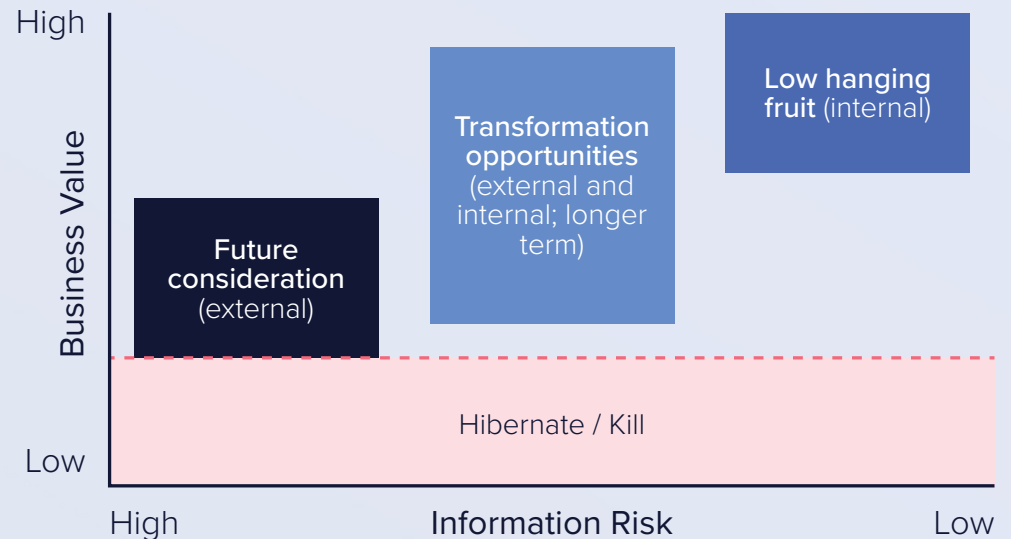


Prioritization strategy

Many firms are prioritizing internal use cases initially due to their lower risk profile and easier implementation. This “dipping a toe” approach allows organizations to learn from internal cases before moving to external implementations that may pose greater risks.

Common internal applications include:

- Search and retrieval of corporate documents and policies
- Automated regulatory change management
- Horizon scanning of market and competitive activities
- First-pass contract and third-party document review
- Large data and online meeting summarization
- Fraud detection





For client-facing applications, there's a priority in making sure the risks are fully understood and a prudence about whether it's the right time to adopt these things.

Amy Longo, Partner, Ropes & Gray LLP

Knowing the limits of technology

Generative AI will remain over-hyped for the foreseeable future. Regulators have already signaled their intent to focus on false or misleading claims over the use of AI ("[AI Washing](#)"). Firms need to exercise care to invest in generative AI approaches that have been thoroughly vetted for specific use cases. Many generative AI approaches will never be suitable for regulated firms, and a separation of those that can be characterized as 'regulatory grade' will eventually occur. Close collaboration between data science teams and those business and compliance stakeholders will continue to be imperative.



When assessing whether and how to incorporate generative AI into business processes, consideration should be given by compliance professionals to the limits of the technology to ensure clarity around how it will be used and for what purposes. Transparency and explainability will be key requirements.

Nina Bryant, Senior Managing Director, FTI Consulting

Cultural transformation

There's a growing recognition that business units need to view data as a strategic asset and that generative AI initiatives should be aligned with clear business outcomes and value propositions.

At the [Reuters AI Momentum event](#), Teresa Heitsenrether, Chief Data and Analytics Officer at J.P. Morgan, said, “Putting generative AI in user hands is like a thousand flowers blooming. You’re seeing the same problem being solved multiple times, now trying to identify common applications. Documents and asking questions of data are targets, but getting businesspeople to think of data as an asset is a cultural shift. Emphasis needs to shift to expected outcomes with a business value defined.”

Any analysis of generative AI-enabled use cases needs to consider the impact on staffing. While [many studies](#) project a transformational impact on workforces, the ability to move staff away from routine, data processing-intensive tasks will not be fast or easy.

Emerging Best Practices

- 1. Establish AI governance councils:** Create dedicated bodies to oversee generative AI initiatives, ensuring alignment with organizational strategies, regulatory compliance, and ethical standards.
- 2. Develop comprehensive evaluation frameworks:** Create detailed checklists and assessment tools that cover all aspects of generative AI implementation, including data protection, cybersecurity, regulatory compliance, and ethical considerations.
- 3. Engage diverse stakeholders:** Involve representatives from various departments, including compliance, legal, IT, and business units, and external advisers in the evaluation process to ensure a holistic assessment of benefits and risks.
- 4. Be aware of technology limits:** Business and compliance teams should remain in constant contact with data science teams to surface false or misleading vendor claims about its technology. Hire simultaneous translation services as needed.



How are stakeholder perspectives integrated into generative AI governance and risk management practices?

Generative AI can be a shiny new toy to some; however, the financial services industry recognizes the importance of balancing innovation with risk mitigation for generative AI use cases.



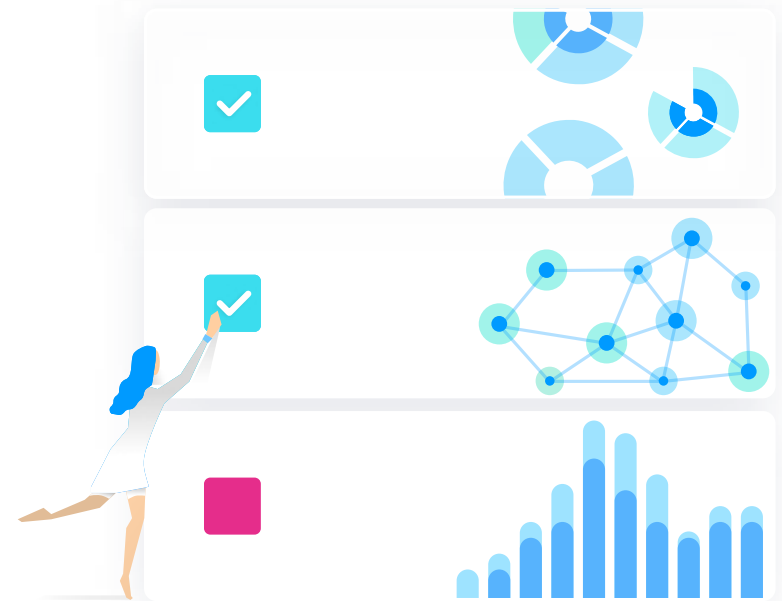
What I'm seeing is a lot of focus on the process up front, and a real effort to try to balance the desire to innovate with the desire to mitigate risk.

Amy Longo, Partner, Ropes & Gray LLP

Generative AI has united functional stakeholders around one common element: the intellectual capital and risk associated with the firm's information. Generative AI can be embedded in, on, around, or with the firm's IP, which has broadened interest in the topic beyond the risk and data science teams.

However, as best practices continue to emerge, there is a notable lack of consistency in how this alignment happens across organizations. Some firms employ sophisticated, inclusive strategies, while others rely on traditional risk management approaches or avoid addressing the issue altogether.

Many organizations rely heavily on external expertise, indicating a shortage of in-house knowledge. This expertise gap underscores the need for substantial internal capacity building in AI governance. Firms are increasingly recognizing the value of diverse stakeholder input in generative AI decision-making processes, aiming to ensure that their strategies are both innovative and responsible.



Emerging Best Practices

- 1. C-level executive risk-aware innovation strategy:** A critical tone-from-the-top agenda item is conveying the objective of maintaining a careful balance between leveraging generative AI's innovative potential and mitigating associated risks.
- 2. Balanced evaluation framework:** Develop sophisticated processes that equally emphasize efficiency and thoroughness in evaluating generative AI use cases. Establish vigorous upfront procedures that enable comprehensive assessment without stifling innovation or creating undue burdens.
- 3. Functional cross-pollination:** Enable users and risk stakeholders to learn from the experiences of other teams, sharing results of early experiments and documenting lessons learned in risk identification and mitigation.
- 4. Internal expertise development:** Invest in comprehensive AI training programs, recruit specialized talent, and cultivate a culture of continuous learning to build robust in-house AI governance capabilities.
- 5. Strategic external partnerships:** Actively engage with industry groups, academic institutions, and peer organizations to stay abreast of best practices.



The background is a dark blue field filled with numerous vertical lines of varying heights. Each line is topped with a small, glowing blue dot, creating a starry or data-point effect. The lines themselves are thin and light blue, and they appear to be slightly curved or wavy, giving a sense of depth and movement. The overall aesthetic is clean, modern, and tech-oriented.

Chapter 2

Regulatory and Risk Implications: How to Be Ready

What methods are being used to identify, assess, and prioritize generative AI risks?

Financial services firms employ various methods that combine traditional risk management frameworks with emerging techniques specifically tailored to address the unique challenges posed by generative AI.

Many organizations are starting with their established technology risk assessment processes as a foundation. These existing frameworks are generally effective for evaluating traditional risks such as security and stability. However, there's a growing recognition that generative AI presents distinct challenges that require adaptation of these processes.

Firms are emphasizing comprehensive data lifecycle management within generative AI systems. This includes rigorous examination of data privacy and security protocols, scrutiny of AI model training processes and data sources, and careful consideration of how proprietary data is used and stored. These measures are crucial for maintaining regulatory compliance, protecting sensitive financial information, and mitigating potential biases that could arise from training data.

Risk assessment procedures are being expanded to include considerations specific to generative AI applications. This includes evaluating system redundancy, disaster recovery capabilities, and business continuity plans in the context of AI-driven systems. Organizations are also developing frameworks to assess risks unique to or enhanced by generative AI, such as output reliance, hallucinations, intellectual property concerns, and the potential for malicious behavior.



A critical component of the risk management strategy is the integration of human oversight throughout the AI lifecycle. Firms are prioritizing extensive human-led testing to ensure the accuracy, reliability, and quality of AI-generated outputs. This human-in-the-loop approach is seen as essential for preventing issues such as AI hallucinations, maintaining the integrity of AI-generated content, and directly addressing explainability concerns in response to regulatory inquiries about system design and decision making.

Emerging Best Practices

- 1. Generative AI risk assessment framework:** Develop a comprehensive framework that acknowledges unique AI risks and adapts existing processes. Regularly update to keep pace with evolving generative AI technologies and industry practices.
- 2. Data lifecycle management protocols:** Employ rigorous protocols focusing on privacy, security, and regulatory compliance.
- 3. Human oversight:** Integrate human oversight throughout the AI lifecycle to balance automation with expert judgment.
- 4. “Trustworthy AI” framework:** Adopt a framework that incorporates ethical considerations alongside technical and operational risks.



While AI can be a boon to an organization, a systematic approach should be taken in its implementation, ensuring that certain guardrails are in place and the AI models and generated work product are continuously validated and enhanced.

Nina Bryant, Senior Managing Director, FTI Consulting

How are firms assessing the impact of generative AI upon specific regulatory obligations?

Financial services firms are actively examining existing regulations and closely monitoring proposed generative AI rules across multiple jurisdictions. They are also watching for enforcement actions that can offer insight into how regulators are defining “explainability” requirements and how they will assess whether generative AI-enabled applications are “reasonably well designed.” Scanning the environment for these events will continue to be an ongoing top priority for most firms.



Whether we’re discussing AI or any other innovation, new technologies often present opportunities for better functioning in more efficient markets. But unfortunately, they can also present opportunities for fraud as well as risks for customers, regulated entities, and the economy at large.

Summer K. Mersinger, Commissioner, CFTC

However, there’s an ongoing debate about the need for new AI-specific regulations. Industry advocacy groups like the Securities Industry and Financial Markets Association (SIFMA) argue that existing regulations are sufficient to encompass AI technologies. They contend that current frameworks, when properly applied, can effectively address the risks and challenges posed by generative AI without the need for additional regulatory burdens.

While recent enforcement actions primarily address basic issues of truthful representation, the industry anticipates more complex cases in the future. These potential cases may delve deeper into the actual operation of AI technologies and their alignment with existing regulations, such as investment

You Should Know

FINRA Regulatory Notice 24-09 reminds firms that FINRA’s rules, which are technology-neutral, continue to apply to the use of AI and generative AI tools. Firms must ensure their use of these technologies complies with existing regulatory obligations (e.g., supervision, communications with the public, books and records).

advisers' fiduciary duty or the best interest rule for broker-dealers. At the most fundamental level, firms can expect regulators to examine if AI-enabled systems are reasonably designed and will expect that firms can defend the methods used by the system to arrive at decisions.



Everyone may be talking about AI, but when it comes to investment advisers, broker-dealers and public companies, they should make sure what they say to investors is true.

Gary Gensler, Chair, U.S. Securities and Exchange Commission

The regulatory landscape is further complicated by varying jurisdictional approaches to AI regulation. Firms are diligently tracking both current and proposed regulations across multiple regions, aware of potential conflicts between regulatory frameworks in different areas, such as Europe and the US.

Key regulatory themes emerging globally include:

- Transparency in disclosures and investor communications
- Explainability in AI decision-making
- Ensuring fairness by eliminating bias and preventing discrimination
- Maintaining human accountability across the AI lifecycle
- Ensuring AI safety and resilience, including protection from cyber threats



Emerging Best Practices

1. **Comprehensive regulatory monitoring:** Track existing rules and proposed regulations across multiple jurisdictions.
2. **Outcome-focused use cases:** Emphasize output and outcomes of generative AI use cases, and not the underlying technologies.
3. **AI with a human component intact:** Leverage AI to enhance compliance with escalating regulatory demands while maintaining human oversight and accountability.
4. **Regulatory scrutiny preparedness:** Establish robust AI risk management programs in anticipation of increased regulatory scrutiny.
5. **Evolving governance structures:** Regularly review and update AI governance structures to ensure they remain effective as the technology and regulatory landscape evolve.

How are firms monitoring developments related to industry standards?

Many firms are utilizing traditional methods like closely following regulatory communications, including consultation papers, webinars, and other published content from regulatory bodies. They are also leveraging industry expertise by relying on specialists who summarize and interpret regulatory statements to provide deeper insights.

However, in spite of the leadership of [NIST](#) in the US and the [EU AI Act](#), forward-thinking firms recognize that relying solely on current regulatory guidance is insufficient. These companies are adopting more proactive approaches to stay ahead of emerging trends, such as monitoring communications from AI development companies to anticipate future technological advancements.




What you need to do is look ahead and recognize not just what AI is today, but where it might be tomorrow, because relying on regulator guidance alone may be insufficient.

Christian Hunt, Founder, Human Risk Limited

Emerging Best Practices

- 1. Comprehensive monitoring:** Compile and analyze up-to-date information on AI-related regulations and standards from diverse sources.
- 2. Cross-disciplinary collaboration:** Work with external experts and participate in industry forums to interpret and apply emerging AI standards.
- 3. Regular internal review:** Systematically assess AI systems and practices against evolving standards, ensuring ongoing compliance and identifying potential impacts on existing practices.
- 4. Future-minded regulatory view:** Look beyond current regulations by monitoring communications from AI development companies to anticipate future technological advancements and their potential implications for industry standards.
- 5. Agile regulatory response:** Track global regulatory differences, especially for firms operating internationally, to navigate varying regulatory landscapes across jurisdictions.



Chapter 3

Generative AI and the Impact on Compliance

How are risks being translated into actionable policies for managing generative AI risks?

Financial services firms are adopting comprehensive approaches to translate legal, regulatory, and IP risks into actionable policies and processes for managing generative AI risks. This involves a two-pronged strategy: developing policies tailored to specific generative AI use cases and updating existing policies to incorporate generative AI considerations across areas like vendor management, privacy, and business continuity.

A focus for firms today is to understand the output of each of the targeted generative AI use cases to recognize where a regulatory or internal policy obligation exists. Is the output of the use case accessible externally, or will it be used to enable decision-making about a product or service of the firm? Or is it accessible only to a firm employee as a productivity tool? Does the output represent value or risk to the firm's business?

An emerging area of focus in policy development is addressing the challenge of "shadow AI," which is the unofficial use of AI tools by employees. Despite formal policies, individuals may utilize generative AI tools for various reasons, including curiosity, productivity enhancement, or fear of job obsolescence. To address this, firms are developing comprehensive strategies that go beyond simple prohibition, which may include:

- Creating clear guidelines for acceptable use of AI tools
- Implementing training and awareness programs
- Developing sanctioned alternatives to popular external AI tools
- Enhancing monitoring capabilities to detect unauthorized AI tool usage
- Establishing open communication channels for employees to discuss AI tool usage
- Regularly updating policies to account for new AI tools and use cases



Emerging Best Practices

- 1. Dual approach to policy updates:** Develop specific generative AI policies while also updating existing policies to incorporate generative AI considerations.
- 2. Rigorous generative AI tool vetting:** Implement a systematic process for evaluating, approving, and monitoring generative AI technologies, involving key stakeholders.
- 3. Balanced innovation and control:** Develop strategies that encourage responsible AI innovation while maintaining effective risk management controls.
- 4. Shadow AI management:** Implement comprehensive strategies to address unofficial AI tool usage, balancing risk mitigation with the potential for innovation.
- 5. Dynamic policy updates:** Regularly review and update AI governance policies to keep pace with evolving technologies and use cases.



Chapter 4

Governance, Accountability and Model Safety

How are organizations implementing generative AI governance?

More firms are integrating generative AI governance by updating crucial policies and processes. These updates often include revising privacy assessments, acceptable use policies, access controls, data retention policies, and third-party risk management evaluations to address generative AI-specific concerns.

Some organizations are looking to address the unique risks associated with generative AI by implementing specialized governance processes. This includes creating new roles and teams dedicated to generative AI oversight and management (“Generative AI Czar”), reflecting an investment in specialized expertise to navigate the complex terrain of AI technologies.




Without embracing and understanding GenAI, compliance officers cannot be very effective at understanding the risks that your business is running

Christian Hunt, Founder, Human Risk Limited

Emerging Best Practices

- 1. Specialized AI governance roles:** Create dedicated positions or teams for generative AI oversight and management.
- 2. Human-AI collaboration framework:** Implement robust processes for human review and judgment in deploying and operating generative AI tools.
- 3. Cross-functional collaboration:** Ensure close collaboration between legal, compliance, IT, and business departments in developing and implementing generative AI governance.
- 4. Ethical AI use guidelines:** Develop and communicate clear guidelines for ethical AI use, including specific boundaries and explanations for these parameters.
- 5. Address the skills gap:** Actively invest in developing AI and generative AI expertise within compliance teams. This could involve training existing staff, hiring professionals with technical backgrounds, or partnering with external experts. Understanding the technology is crucial for effective risk management and regulatory compliance.



How are firms providing due diligence on existing applications that are now embedding generative AI?


The integration of Copilot into Microsoft Teams has caused more firms to recognize the need for new approaches to due diligence for existing applications that are now embedding generative AI. Rather than relying solely on initial approval processes, firms are implementing ongoing monitoring systems to reassess the risk profiles of these evolving applications.



You could have brought something in on the presumption it was one thing, and it becomes something fundamentally different. The approval process that would have got that in through the door wouldn't have asked the kinds of questions that the addition of, say Copilot embedded in Microsoft would add into it.


Christian Hunt, Founder, Human Risk Limited

Firms are paying close attention to software update schedules and release notes, acknowledging that AI capabilities can be introduced at any time, potentially altering the risk landscape. Firms are also emphasizing employee awareness and engagement, encouraging an “if you see something, say something” culture.



Emerging Best Practices

- 1. Continuous due diligence:** Avoid risk with applications integrating generative AI. Move beyond one-time approval processes to ongoing risk assessment, including processes for documenting and tracking changes in generative AI capabilities within existing applications over time.
- 2. Adaption to risk profile changes:** Anticipate and prepare for potential risk profile changes resulting from generative AI integration rather than reacting to changes after they occur.
- 3. Ongoing monitoring:** Establish systems for ongoing monitoring and reassessment of risk profiles for applications integrating generative AI features.
- 4. Controlled update management process:** Adopt a process where updates to generative AI-integrated applications are limited in use until independently evaluated and approved.
- 5. Awareness and reporting:** Foster a culture of awareness and reporting among employees regarding generative AI integrations and potential issues.



How are firms evaluating and selecting specific generative AI models?

Increasingly, more firms appear to be moving toward a platform-agnostic approach to model selection to mitigate the concentration of risk relying upon a single provider's foundational model offerings. Instead, they're developing and applying specific criteria to specific use cases considering their expected benefit and risk profile.



I would expect that the trend of company-specific, industry-specific GPTs and models that are easier to use and cheaper to run, will allow control teams to apply domain expertise against a specific use case. You use the right tool for the right job.

Matthew Bernstein, Information Governance Strategist, MC Bernstein Data

There's growing recognition within the industry that optimal results often come from integrating and customizing multiple AI tools rather than relying on a single solution. This illustrates the importance of adopting NIST AI standards across all segments of financial services. Firms are conducting thorough performance testing of various model combinations to ensure cost-effective solutions tailored to their organizational needs.

Financial institutions are also developing sophisticated approaches to evaluate and select generative AI models, focusing on comprehensive assessment frameworks and flexible integration strategies. These assessment frameworks typically consider multiple factors:


- Defining the business problem or opportunity the AI model aims to address
- Assessing potential revenue growth, cost savings, efficiency gains, and overall impact on business objectives
- Evaluating potential regulatory, reputational, operational, and financial risks
- Defining and measuring key indicators of model effectiveness and efficiency specific to the use case
- Examining potential bias, fairness, transparency, and societal impacts of the model
- Analyzing technical feasibility, required resources, and integration challenges
- Assessing the availability, quality, and accessibility of necessary data and supporting technology
- Considering the model's ability to grow with the business and remain relevant over time

Companies are increasingly recognizing the need for subject matter experts in compliance and audit roles who thoroughly understand generative AI capabilities. These experts are crucial for designing effective governance frameworks and conducting meaningful risk assessments.



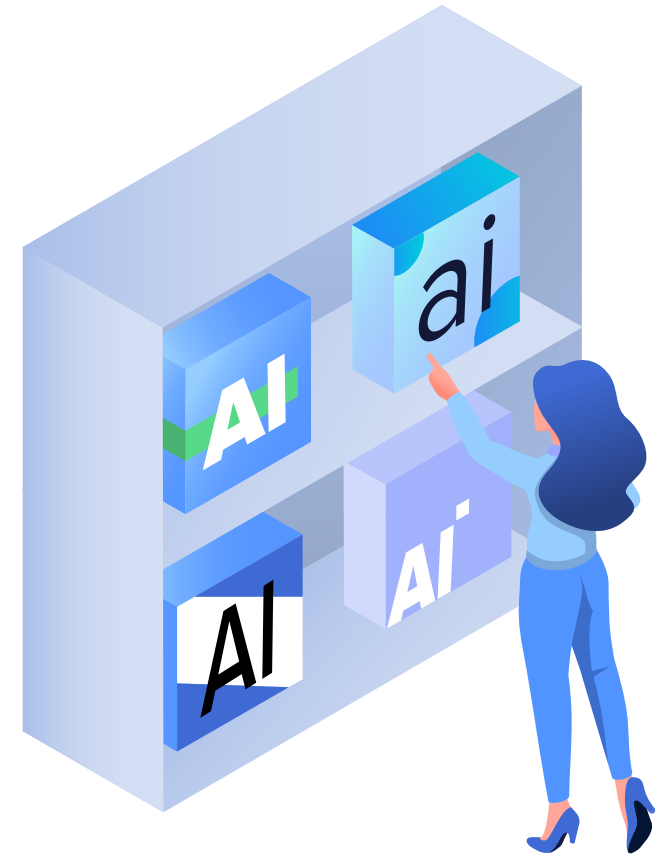
If those managing generative AI integration aren't thoroughly researching and understanding its capabilities, they're demonstrating a lack of subject matter expertise and failing to recognize the near certainty of unofficial use within the organization.

Christian Hunt, Founder, Human Risk Limited



Emerging Best Practices

- 1. Platform-agnostic approach:** Develop performance metrics specific to intended use cases rather than defaulting to a single provider's offerings, allowing for more tailored and cost-effective solutions.
- 2. Off-channel policies for generative AI tools:** Adopt policies to address the use of these tools by employees on personal devices or external platforms. This could involve creating sanctioned alternatives, implementing more comprehensive monitoring, or fostering a culture of open communication about AI tool usage.
- 3. Generative AI expertise:** Cultivate expertise in generative AI by fostering technological awareness across traditionally non-tech roles and recruit or train subject matter experts who thoroughly understand generative AI capabilities.
- 4. Comprehensive model selection criteria:** Apply model criteria that balance performance, business value, risk mitigation, ethical considerations, and implementation feasibility.
- 5. Data protection and control:** Make data security a priority when evaluating generative AI models, particularly for sensitive industries like financial services.



Conclusion

The entry of generative AI into the financial sector marks a significant turning point in how firms operate, make critical decisions, and engage with the market. Amidst the swirl of excitement, the dual narratives of potential and peril are unmistakably present.

This transformative journey is not without its challenges—from regulatory complexities and operational risks to ethical considerations. However, the overarching sentiment is one of cautious optimism, emphasizing a balanced approach to harnessing the power of generative AI while navigating its multifaceted challenges.

Financial institutions, guided by insightful leadership and bolstered by the expertise of compliance officers and risk stakeholders, are carving paths through this uncharted territory. The collaborative spirit seen across business, technological, legal, and human risk disciplines underscore a collective commitment to not only mitigating risks but also unleashing the transformative potential of generative AI.





Smarsh®, the global leader in communications data and intelligence, enables companies to transform oversight into foresight by surfacing business-critical signals in all their digital communications. Regulated organizations of all sizes rely upon the Smarsh portfolio of cloud-native digital communications capture, retention and oversight solutions to help them identify regulatory and reputational risks within their communications data before those risks become fines or headlines.

Smarsh serves a global client base spanning the top banks in North America, Europe and Asia, along with leading brokerage firms, insurers, and registered investment advisers and U.S. state and local government agencies. To discover more about the future of communications capture, archiving and oversight, visit www.smarsh.com

Smarsh provides marketing materials for informational purposes only. Smarsh does not provide legal advice or opinions. You must consult your attorney regarding your compliance with applicable laws and regulations.

E-Book - 09/24



1-866-762-7741



www.smarsh.com



@SmarshInc



SmarshInc



Company/smarsh