BERMUDA MONETARY AUTHORITY

Wednesday, 30 July, 2025

DISCUSSION PAPER

The Responsible Use of Artificial Intelligence in Bermuda's Financial Services Sector

Submissions to be received by Tuesday, 30 September 2025

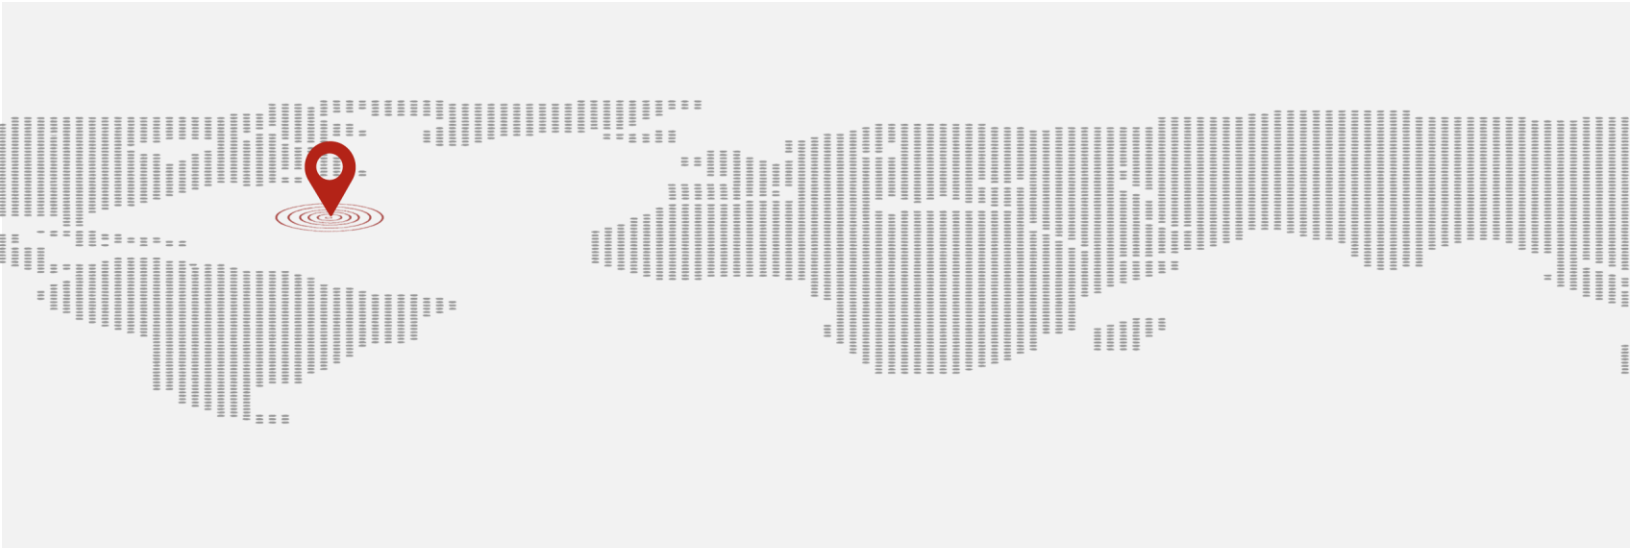

I. Table of Contents

Industry stakeholders and other interested persons are invited to provide feedback on the proposals set out in this Discussion Paper (DP) by 30 September 2025. There are two options for submitting feedback:

1. Respond to the questions and provide comments using the survey form located here: https://forms.office.com/r/1wWTYf96wX

2. Alternatively, you can choose to email your responses to policy@bma.bm . For this option, please include your company's name, followed by 'Discussion Paper – *The Responsible Use of Artificial Intelligence in Bermuda's Financial Services Sector'* in the subject line of the email.

II. Introduction

1. The financial services landscape is undergoing a profound transformation driven by the rapid advancement of Artificial Intelligence (AI) technologies. As Bermuda's financial services regulator, the Bermuda Monetary Authority (Authority or BMA) recognises the need to understand and supervise how regulated firms deploy AI-enabled solutions to maintain Bermuda's reputation as a premier financial services jurisdiction.

2. This DP builds upon the BMA's previous work in the insurance sector in 2022[1] and reaffirms its commitment to fostering innovation while promoting financial stability, customer protection and the integrity of Bermuda's financial system.

3. The BMA has long been recognised for having an adaptive and innovative approach to regulation that can successfully balance innovation with appropriate regulatory oversight. Within this DP, the BMA seeks to extend a forward-thinking approach to the rapidly evolving domain of AI across all of Bermuda's financial services sectors. The Authority aims to address the unique challenges and opportunities that AI presents while fostering an environment where responsible innovation can flourish.

4. This DP covers:

    a) An overview of AI applications within Bermuda's financial services;

    b) A general assessment of global regulatory approaches with respect to AI;

    c) An identification of sector-specific risks and opportunities presented by AI adoption; and

    d) An outline of potential pathways for developing an appropriate regulatory framework.

5. This DP seeks to:

    a) Foster responsible innovation in AI by offering regulatory clarity and supportive engagement that can allow new solutions to be developed safely;

    b) Ensure Bermuda's regulatory framework remains fit-for-purpose while adhering to international standards; and

    c) Solicit targeted feedback from industry participants to inform the Authority's future guidelines on responsible AI use.

6. This collaborative approach reflects the BMA's commitment to maintaining open dialogue with industry stakeholders to navigate the complexities of the collective integration of AI within financial services.

7. Through this discussion process, the BMA seeks to develop an outcomes-based regulatory framework for AI governance and oversight that leads to responsible AI adoption. For indicative purposes, the following steps are summarised in the table below.

Phase	**Target Date**
Consultation closes	30 Sept 2025
Analyse feedback	Q4-2025
Follow-up consultations and industry workshops	Q1-2026
Final proposal	Q3-2026

[1] Bermuda Monetary Authority (2022), *Bermuda Insurance Sector Artificial Intelligence and Machine Learning Survey – 2022 Report*.

III. Executive Summary

8. The BMA recognises AI as a transformative force in financial services. When properly governed, AI offers significant opportunities to enhance efficiency, improve customer outcomes and drive innovation. However, this must be balanced against managing the complex risks inherent in Bermuda's distinctive ecosystem of sophisticated institutional clients and cross-border arrangements. Rather than stifling technological advancement, the Authority seeks to create a regulatory environment that enables financial institutions to harness AI's potential, while maintaining the high standards of stability and integrity that define Bermuda's reputation as a premier financial jurisdiction.

9. Building on the BMA's 2022 insurance sector survey[2], the Authority acknowledges the dramatic acceleration of AI adoption globally, particularly in the context of generative AI systems. This acceleration necessitates a more proactive regulatory stance, as regulated financial institutions are increasingly implementing advanced AI systems across core business functions. For example, catastrophe modelling and underwriting, as well as compliance monitoring and customer service, often lack comprehensive governance frameworks.

10. This DP reviews global regulatory developments from international standard-setting bodies, including the following:

    - International Association of Insurance Supervisors (IAIS)
    - Financial Action Task Force (FATF)
    - International Monetary Fund (IMF)
    - International Organization of Securities Commissions (IOSCO)
    - Bank of International Settlements (BIS)

    This DP will also review major legislative initiatives in key jurisdictions, including:

    - The EU Artificial Intelligence Act (EU AI Act)
    - The UK's innovation framework
    - Singapore's practical risk management guidance

- US National Institute of Science and Technology (NIST) voluntary framework

The BMA's analysis of these various initiatives reveals convergence around key principles despite varying implementation approaches, namely:

- An emphasis on governance and board of director's (board) accountability
- Risk-proportionate oversight
- Requirements for human oversight in critical applications
- Transparency and explainability obligations
- Integration with existing regulatory frameworks rather than creating separate AI regimes

[2] Bermuda Monetary Authority (2022), *Bermuda Insurance Sector Artificial Intelligence and Machine Learning Survey - 2022 Report*.

These global developments provide essential insights for Bermuda's approach while recognising the jurisdiction's unique position as a sophisticated international financial centre serving primarily institutional clients.

11. Based on this analysis, the BMA proposes an outcomes-based risk management framework emphasising governance and oversight, with ultimate accountability resting with an entity's board. The framework addresses critical components, including AI identification and inventory management, alongside comprehensive risk assessment across five key dimensions: impact severity, autonomy and human oversight, complexity and explainability, data sensitivity, and deployment context and scale. The proposed framework then outlines governance and risk considerations for financial institutions regarding:

    - Robust data management
    - Model development and validation
    - Human oversight requirements
    - Explainability and fairness considerations
    - Ongoing monitoring and change management
    - Third-party risk management
    - Generative AI-specific controls
    - Cybersecurity and operational resilience measures

    The proposed approach attempts to address both individual institutional risks and market-wide systemic considerations, including concentration risks from common AI providers and potential cascading failures across interconnected financial systems.

12. Acknowledging Bermuda's diverse financial marketplace and varying institutional capabilities, the framework applies the proportionality principle to scale governance requirements based on a financial institution's business profile, size, complexity, AI maturity and the nature of customer relationships (i.e. retail versus institutional). This principles-based, risk-proportionate approach balances innovation with the appropriate safeguards while providing practical implementation guidance, including phased approaches, collaborative strategies for smaller financial institutions, and resource development requirements.

IV. Background on the Use of AI for Financial Services

13. The global financial services industry is experiencing accelerated AI adoption, propelled by several interconnected forces. Technological breakthroughs in computing power, data availability and AI algorithms have dramatically expanded solution capabilities while reducing costs. Simultaneously, the digital transformation of customer services has heightened expectations that AI will deliver round-the-clock availability, personalisation and seamless experiences. Traditional financial institutions now face significant competition from FinTech startups and 'tech giants' that leverage AI, forcing established players to accelerate their own AI adoption strategies.

14. Recent market data from CB Insights[3] illustrates the accelerating momentum of AI adoption globally. In Q1 2025, AI funding surged to a record $66.6 billion—a 51% quarter-over-quarter increase—with the median deal size reaching a four-year high of $5 million, indicating growing confidence in established market leaders. Meanwhile, consolidation among enterprise AI companies foreshadows integration possibilities for

[3] CB Insights (2025), *State of AI Q1'25 Report*.

banking, payments and digital asset businesses, where unified AI platforms could transform cross-border transactions, compliance monitoring and customer engagement across Bermuda's sophisticated financial marketplace.

15. AI presents significant opportunities for process automation, error reduction and resource optimisation. The growing complexity of regulatory requirements has also driven interest in AI-powered compliance solutions that adapt to changing rules and monitor transactions more effectively. Additionally, the explosion of available data creates both challenges and opportunities, with AI tools capable of extracting actionable insights from previously unmanageable information volumes.

16. As observed by the Authority from its initial work in this area in 2022, the industry appears to be transitioning from pilot AI deployments to integration within core business functions. This AI adoption is expanding from large financial institutions to smaller entities through partnerships and third-party providers, making the adoption of AI a critical consideration for regulators across the market. Within Bermuda's sophisticated financial environment, the BMA expects AI systems to be deployed across various contexts, with different levels of risk and varying use cases.

17. From the outset, the BMA recognises that AI exists on a spectrum of sophistication and autonomy, from rule-based systems to advanced machine learning models capable of adapting and improving without explicit programming. This DP primarily focuses on modern AI applications that involve significant degrees of automation, learning capabilities, and decision-making or decision-support functions.

Definition and Types of AI

18. For the purposes of this DP, AI systems are technology solutions that can analyse data, learn from patterns, and make decisions or recommendations with varying degrees (or the absence) of human oversight.

19. In the broader financial services context, AI encompasses systems that exhibit two critical characteristics: autonomy (the ability to operate and make decisions with limited human intervention) and adaptiveness (the capacity to improve results or adjust behaviour based on new information or feedback). This definition can include traditional machine learning models, generative and agentic AI applications, and automated decision-making systems, but excludes simple rule-based programmes or static mathematical models that require explicit programming for each decision scenario.

20. Financial institutions should recognise that systems not formally labelled as ‘AI’ may still warrant AI governance considerations if they demonstrate these autonomous and adaptive characteristics in their operations.

The Bermuda Context

21. Bermuda has established itself as a premier global financial centre specialising in insurance, reinsurance, banking, investment management and digital assets, serving institutional clients and retail clients often through cross-border arrangements.

22. Wholesale institutional markets leverage AI for complex risk modelling, sensitive commercial data processing, and high-stakes decision-making that affects global markets. This requires nuanced regulatory approaches that balance innovation with appropriate oversight.

23. In the retail markets, particularly in sectors where customers may lack the ability to scrutinise complex model outputs, an institution’s policy should promote fairness and transparency through plain-language disclosures, bias testing and clear customer support mechanisms. These measures help build trust and ensure that affected

customers have accessible pathways to raise concerns or seek clarification where AI-driven decisions materially impact them.

V. BMA's Early Initiatives

24. The BMA's 2022 insurance sector AI survey revealed a 38% adoption rate, with primary applications in cybersecurity, underwriting and claims management. Key challenges identified included a lack of written strategies, expertise gaps and concerns about explainability and auditability.

25. Since 2022, Generative AI has become mainstream, necessitating updated regulatory approaches addressing model explainability, bias detection and governance frameworks for large language models and other generative systems.

Cross-Border, Cross-Sectoral Implications

26. The Authority further recognises that AI adoption represents a cross-border and cross-sectoral phenomenon, which requires a holistic regulatory approach across the financial services sectors it regulates. AI solutions frequently transcend traditional industry boundaries, with similar technologies deployed across banking, insurance, investments and other financial sectors. To add to this, many AI systems used by Bermuda-based entities are developed, maintained or operated by international technology providers, creating complex jurisdictional considerations.

27. AI systems typically involve cross-border data transfers, raising important questions about data sovereignty, privacy and compliance with multiple regulatory regimes. As financial institutions increasingly rely on common AI platforms or methodologies, systemic risks emerge that transcend traditional sectoral boundaries. Their interconnected nature has prompted the BMA to consider a coordinated regulatory approach that balances sector-specific considerations with broader principles that are applicable across the financial services landscape.

The BMA's Business Plan: AI for Enhanced Operations and Supervision

28. In alignment with its commitment to technological innovation, the BMA has incorporated AI enhancement into its own operational and supervisory processes. Over the last five years, the BMA has created and grown a separate Data Science and Artificial Intelligence Department, with the mandate of developing data analytics and AI capabilities to support the identification of emerging risks more effectively and allocate supervisory resources proportionately.

29. These initiatives align directly with the BMA's 2025 Business Plan objectives for technological advancement and supervisory innovation. As outlined in the Business Plan, the Authority has committed to leveraging technology, including AI, as part of its IT Strategy 2030, positioning the Authority to not only regulate AI effectively, but also to benefit from its transformative potential in risk-based supervision and regulatory efficiency. This practical experience also provides the BMA with valuable insights that directly inform its approach to regulation and supervision, ensuring that future guidance is both practical and forward-looking.

VI. AI Applications in Financial Services

30. AI applications across Bermuda's financial sectors offer diverse and transformative benefits. They can significantly enhance operational efficiency, automate complex processes, strengthen compliance and risk management, and improve the speed and accuracy of decision-making.

31. In the insurance sector, AI could significantly enhance underwriting precision for catastrophe modelling and emerging risks. Generative AI supports claims processing automation and fraud detection in large commercial policies. For reinsurance, AI can analyse vast datasets of global risk patterns to optimise pricing

for excess-of-loss treaties. These capabilities are particularly valuable for complex international risk transfer mechanisms and alternative risk financing structures. For example, the same models could underpin parametric micro-insurance for climate-exposed households and small and medium enterprises (SMEs) , allowing carriers to price and settle claims significantly more rapidly, thereby helping to close part of the global protection gap.

32. For Bermuda's banking sector, which serves both international businesses and local residents, AI can enable sophisticated Know Your Customer (KYC) and Anti-Money Laundering (AML) compliance solutions tailored to various clientele. Generative AI can power personalised wealth management services and investment recommendations for institutional clients. In the domestic market, AI chatbots can provide 24/7 customer service, while machine learning enhances credit decisions and fraud monitoring for local retail customers. As an example, credit-scoring models that ingest cash flow and utilities data can potentially broaden lending to a segment of the retail customers without relaxing underwriting standards or increasing risks. These tools may assist banks in maintaining their competitive edge while navigating complex global compliance requirements.

33. Bermuda's trust and corporate service providers may leverage generative AI to streamline complex document preparation and compliance verification for international corporate structures. AI can enhance due diligence processes by automating background checks and ongoing monitoring of high-net-worth individuals and corporate entities. Machine learning models can analyse beneficial ownership structures and cross-border transactions to identify potential risks. Other use cases include handling the increasingly complex and cross-border regulatory and tax reporting requirements of the firms and the structures under management. Additionally, there is a growing need to manage the increasingly complex family relationships administered on a cross-border/multigenerational basis, which requires specialist advice from service providers in Bermuda.

34. These capabilities could greatly benefit Bermuda's sophisticated market, where providers manage complex asset structures and international holding companies that require meticulous governance.

35. For Bermuda's investment fund sector, AI algorithms optimise portfolio allocations for insurance-linked securities and alternative investments. Generative AI has the capability to produce more advanced, detailed and real time market analyses and regulatory reporting documentation. Machine learning models can analyse global market data to predict trends specific to Bermuda's specialised investment vehicles. AI-powered governance tools can also monitor fund compliance with international standards at a much faster rate than the current practices. These applications are especially valuable given Bermuda's fund industry, which exceeds $290 billion and caters to institutional investors seeking sophisticated offshore investment strategies.

36. In Bermuda's Money Service Business (MSB) sector, AI could enhance cross-border payment efficiency through real-time processing and settlement, fraud detection and currency optimisation algorithms. Generative AI can also power multilingual customer service support systems for diverse clientele. Machine learning models can identify suspicious transaction patterns specific to Bermuda's unique geographic position in a matter of minutes, instead of days. These applications could help MSBs maintain competitive pricing while meeting stringent regulatory requirements, particularly important for facilitating remittances for Bermuda's international workforce.

37. For Bermuda's evolving payments landscape, AI could streamline transaction routing and optimise cross-border payments crucial for international business operations. Generative AI may also enhance payment security through advanced behavioural analytics to detect unusual payment patterns. Machine learning also enables real-time currency conversions at optimal rates. For domestic payments, AI can significantly

improve the user experience through seamless authentication. These capabilities support Bermuda's dual role in serving sophisticated international businesses while meeting local payment needs, reinforcing the island's position as a progressive financial hub.

38. Bermuda's pioneering digital asset regulatory framework has attracted innovative businesses that leverage AI for enhanced compliance and security. Generative AI can facilitate automated smart contract creation and regulatory disclosures at a significant speed and with precision. Machine learning models can also monitor blockchain transactions in real-time to detect suspicious patterns and ensure compliance with Bermuda's Digital Asset Business Act (DABA) regulations, and potentially, facilitate automated implementation of regulatory and supervisory tools for DAB entities. AI-powered risk models can analyse cryptocurrency market volatility to protect client assets. Decentralised Large Language Models, while scarce at the moment, may likely achieve scalability and may soon become more accessible to blockchain networks, unlocking further opportunities in this space.

39. Given Bermuda's role as a global financial hub hosting sensitive data for international and domestic customers, AI-powered cybersecurity is essential. Machine learning algorithms can detect anomalous network behaviour, indicating potential breaches before traditional systems would recognise them. Generative AI can simulate sophisticated attack scenarios for testing security infrastructure. AI enhances threat intelligence by analysing global cyber threat patterns relevant to Bermuda's financial sector. Natural language processing has the capability to monitor dark web activities for mentions of Bermuda-based financial institutions. It is worth noting that from the 2022 BMA survey, cybersecurity topped the list in the current use of AI applications in the insurance sector. With the rise of Generative AI, these advanced capabilities will help maintain Bermuda's reputation as a secure jurisdiction for complex financial operations in an increasingly interconnected global landscape.

40. AI can significantly strengthen operational resilience across Bermuda's financial services by predicting potential system failures before they occur. Machine learning models can optimise infrastructure capacity planning during peak processing periods. Generative AI can be used to create and test dynamic business continuity scenarios and stress testing specific to Bermuda's island geographical location and potential hurricane impacts, on an ongoing basis rather than a point-in-time exercise. AI-powered monitoring systems can also continuously assess third-party service provider risks. These capabilities are especially crucial for Bermuda's financial institutions that need to maintain continuity of critical services for international clients despite geographic isolation and natural disaster risks, while meeting regulatory expectations for operational resilience.

41. Across all Bermuda's financial sectors, AI can also significantly enhance AML/KYC and sanctions compliance capabilities. Generative AI automates complex regulatory reporting and creates detailed risk assessments. Machine learning models detect subtle patterns in transactions that might indicate money laundering. Natural language processing scans communications for compliance risks. While the current Bermuda AML and Sanctions frameworks are designed to be technology agnostic, these tools are particularly valuable for the market (and the BMA, given Bermuda's international business focus. They can help financial institutions maintain the highest compliance standards while efficiently managing the extensive documentation requirements of global financial regulations.

42. As a whole, these applications demonstrate how AI can strengthen Bermuda’s position as a leading financial hub by driving innovation, increasing efficiency and enabling new service models across sectors. AI also has the potential to deliver meaningful improvements in customer outcomes by shrinking the global insurance protection gap, lowering the cost of remittances, and expanding fair access to credit and investment opportunities, particularly when applied responsibly and with the appropriate safeguards.

43. However, the advantages of AI must be balanced against the potential risks it introduces. Without robust governance, explainability, oversight and ethical design, AI systems could potentially cause harm to customers, amplify biases or destabilise critical market functions. The BMA's regulatory approach, therefore, aims to ensure that financial institutions remain accountable for the safe and fair use of AI. Innovation must be pursued within a framework that upholds customer protection, financial stability and the integrity of Bermuda's financial system.
44. Further, the BMA aims to design a technology neutral and outcomes-focused regulatory approach, allowing financial institutions the flexibility to innovate while ensuring appropriate risk management. Rather than focusing on prescriptive rules that might quickly become obsolete, the BMA instead emphasises governance frameworks that can adapt to technological evolution while maintaining consistent protection standards. This approach positions Bermuda as a jurisdiction where AI innovation can thrive within a robust regulatory framework.

Discussion Questions for the Industry

To better understand the current landscape and future direction of AI adoption within Bermuda's financial services sector, please answer the following questions:

1. What AI systems or technologies has your organisation already implemented or is currently implementing? Please specify whether these are used for institutional clients, retail customers or internal operations.
2. What AI application(s) is your organisation planning to implement in the next 12-24 months? If applicable, how do you expect these to differ in approach between your institutional and retail client bases?
3. In which business areas or units does your organisation currently implement or plan to implement AI systems?
4. What do you see as the primary barriers or challenges to AI adoption in your specific sector? If applicable, please distinguish between challenges specific to serving institutional versus retail clients.
5. For financial institutions operating globally: How do you currently manage AI governance across multiple jurisdictions? What challenges do you face in ensuring AI systems comply with varying international regulatory requirements while maintaining operational efficiency?
6. How does your approach to AI governance and transparency differ when serving sophisticated institutional clients versus retail customers? What specific considerations apply to Bermuda's predominantly wholesale market context?

VII. Global Initiatives and Legislative Updates

45. The rapid advancement of AI systems has prompted regulatory bodies and governments worldwide to develop comprehensive frameworks ensuring responsible AI use in financial services. International

standard-setting bodies, including the IAIS[4], FATF[5], IMF[6], IOSCO[7] and BIS[8] have established common principles emphasising risk-based, proportionate approaches with consistent themes of board accountability, governance frameworks, third-party oversight, model robustness and fairness considerations. These frameworks universally assign the ultimate responsibility to financial institutions for AI outcomes, regardless of whether systems are developed internally or by external vendors.

46. Regional regulatory approaches demonstrate varying strategies while maintaining core governance principles. The EU's comprehensive AI Act[9] establishes risk-based classifications with significant penalties, while the UK[10] pursues a pro-innovation, principles-based framework through sector-specific regulator guidance. Singapore's MAS[11] has developed practical AI risk management guidance with ongoing industry engagement. In the US, the NIST AI Risk Management Framework[12] provides voluntary and flexible implementation pathways. The UK has also introduced an AI Cyber Security Code[13] to address AI-specific vulnerabilities. The UAE[14] has established a dedicated AI regulatory office with enforcement powers, while Australia[15] emphasises ethics-based approaches through non-binding frameworks.

47. These global developments reveal convergence around several key regulatory principles despite different implementation approaches, namely:

    - Emphasis on governance and Board accountability
    - Risk-proportionate oversight scaling with system impact and complexity
    - Requirements for human oversight in critical applications
    - Transparency and explainability obligations
    - Integration with existing regulatory frameworks rather than creating entirely separate AI regimes

    Most jurisdictions focus on outcomes-based regulation that allows for technological flexibility while ensuring appropriate safeguards, particularly for high-risk applications that affect customer outcomes or financial stability.

48. These frameworks provide valuable insights while recognising Bermuda's unique position as a sophisticated international financial centre. The Authority will consider how varied global approaches can inform a Bermuda-specific model that balances local innovation needs with international alignment, building on its pragmatic oversight in emerging technologies. The evolving nature of these initiatives supports a principles-based, risk-proportionate approach that can adapt as international standards continue to develop, positioning Bermuda as a trusted jurisdiction for responsible AI adoption in financial services.

---

[4] IAIS (2024), *Draft Application Paper on the supervision of artificial intelligence - November 2024* .
[5] FATF (2021), *Opportunities and Challenges of New Technologies for AML/CFT Report - July 2021*.
[6] IMF (2024), *Global Financial Stability Report, Chapter 3: Advances in Artificial Intelligence: Implications for Capital Market* Activities.
[7] The Board of the IOSCO (2021), *The use of artificial intelligence and machine learning by market intermediaries and asset managers – Final Report*.
[8] BIS (2024), *Annual Economic Report June 2024*.
[9] European Union (2024), Artificial Intelligence Act.
[10] UK Government (2023), *Policy Paper: A pro-innovation approach to AI regulation*.
[11] Monetary Authority of Singapore (2024), *Artificial Intelligence Model Risk Management: Observations from a Thematic Review*.
[12] NIST (2023), Artificial Intelligence Risk Management Framework (AI RMF 1.0).
[13] UK Government (2025), *Guidance: Code of Practice for the Cyber Security of AI*.
[14] UAE (2024), UAE's International Stance on Artificial Intelligence Policy.
[15] Australian Government (2019), Australia's AI Ethics Principles.

VIII. Proposed Risk and Governance Considerations and Supervisory Expectations

49. Building on these global regulatory insights, the BMA proposes a tailored approach for Bermuda's unique financial services environment.

50. The global initiatives outlined in the previous section provide a foundational framework for the financial services sector to build upon. For the Bermuda market specifically, these broad principles can be translated into practical risk management approaches that address the unique characteristics of AI systems.

51. The BMA has reviewed key global initiatives and updates on AI and has also considered the unique characteristics of the Bermuda market. With these learnings in mind, the Authority offers the following risk considerations as a starting point to identifying, assessing and mitigating AI-related risks across key domains in financial services.

52. The BMA proposes a comprehensive risk management framework emphasising governance and oversight. This outcomes-based approach focuses on results rather than prescriptive technical implementations, with ultimate accountability resting with the boards of directors. The succeeding sections outline the Authority's proposed governance considerations for financial institutions and its initial supervisory expectations aligned with the perceived risks in each area.

Governance and Oversight

53. A comprehensive AI Governance framework is essential for AI risk management by financial institutions. Without robust governance structures, financial institutions face unclear accountability for AI decisions and outcomes. Effective governance ensures AI systems align with organisational values and regulatory requirements. Poor governance can result in uncontrolled AI deployments, regulatory violations and misalignment between AI system objectives and the financial institution's risk tolerance. Senior leadership may approve AI systems they do not fully understand, creating accountability gaps when AI failures occur.

54. Strong governance requires clear accountability at board and senior management levels, with sufficient AI literacy to challenge decisions and set appropriate risk appetites. Organisations should establish transparent policies, well-defined roles and responsibilities, and proper training programmes.

55. The Three Lines of Defence (3LOD) Model provides a structured framework for AI governance within financial institutions, ensuring multiple checkpoints for identifying and mitigating AI risks before they manifest as operational failures or regulatory breaches. This tiered approach encompasses business units, risk management and compliance functions, and internal audit.

Outcomes-focused Approaches and the Ultimate Responsibility of the Board

56. While the purpose of this DP is to develop a framework that is specifically tailored to the risks and technical challenges of deploying AI systems, the BMA's approach to AI governance is fundamentally outcomes-based, focusing on the desired results to be achieved rather than prescribing specific technical implementations.

57. This approach should provide financial institutions the flexibility to develop AI systems tailored to their unique business contexts while ensuring appropriate safeguards. Central to this framework is the principle that ultimate responsibility and accountability for AI outcomes rest with the board. While the board may delegate day-to-day design, testing and monitoring, it remains accountable for establishing clear risk appetites, ensuring that governance structures include clear reporting lines, and receiving timely assurance that AI systems operate within approved parameters throughout their lifecycle. While this accountability is inalienable, supervision can be exercised through board committees, internal audit reviews and independent

attestations. By emphasising outcomes and board accountability, the BMA aims to foster innovation while ensuring that oversight remains commensurate with risk and that responsibility for AI-driven decisions remains clearly defined at the highest levels of organisational governance.

Risk-Based AI Programmes

58. AI systems pose unique risks requiring structured management approaches to prevent harmful customer outcomes in financial institutions. Without risk-based programmes, organisations may apply inadequate controls to high-risk AI applications or waste resources on excessive measures for low-risk systems, potentially leading to biases, privacy breaches or insecure systems that cause financial and reputational damage.

59. Financial institutions should implement policies, programmes and procedures for responsible AI use, such as that of (but not limited to) NIST's Map-Measure-Manage framework[16]—recognising the context, identifying and assessing risks, and taking appropriate actions. Risk management programmes should cover AI systems across the entire financial services lifecycle.

The Risk Assessment Framework

60. The BMA proposes that financial institutions implement comprehensive risk assessment frameworks that evaluate AI systems across the following five key dimensions:

    a) Impact Severity: The potential consequences of system failure or malfunction on customers, business operations, and the broader financial system, with externally deployed systems generally presenting higher risk profiles than internal systems;

    b) Autonomy and Human Oversight: The degree of human involvement in decision-making and intervention capability, particularly for critical operations, important business services, and regarding interfaces between AI systems and external stakeholders, including the risk that operators acting only as a 'fail-safe' become over-reliant on model outputs, reducing vigilance (known as 'automation bias');

    c) Complexity and Explainability: The transparency, interpretability and explainability of the AI model and its decision-making processes;

    d) Data Sensitivity: The nature and sensitivity of personal and financial information being processed by the system; and

    e) Deployment Context and Scale: The operational environment (internal staff support versus direct customer or market interface) and the scope of deployment, including the number of customers, transactions, or business processes potentially affected.

61. These risks can be demonstrated in a practical context by:

    a) High-Risk Examples: Automated underwriting systems making binding coverage decisions without human review; AI-powered investment allocation systems managing client portfolios; automated loan approval systems for retail customers;

[16] NIST (2023), Artificial Intelligence Risk Management Framework (AI RMF 1.0).

b) Medium-Risk Examples: AI-powered fraud detection systems that flag transactions for human investigation; chatbots providing preliminary customer service with human escalation paths; risk assessment tools supporting human decision-makers; and

c) Low-Risk Examples: Internal portfolio analysis tools used for strategic planning; back-office process automation with human oversight and data analytics supporting management reporting.

The Risk Materiality Assessment

62. Risk management approaches should be proportionate to the assessed risk levels, with enhanced governance for applications that affect critical business functions or customer outcomes. Inadequate risk assessment could lead to catastrophic failures in high-risk applications while wasting resources on excessive controls for low-risk systems. Financial institutions may overlook emerging risks or fail to identify potentially harmful outcomes for customers or the broader financial market.

63. Risk assessment frameworks should include both quantitative and qualitative dimensions to properly allocate management resources proportionally to each AI system's risk. Assessment criteria should evaluate the degree of potential harm to relevant counterparties—the severity of adverse economic impact that counterparties might experience from an unfavourable outcome. Financial institutions should consider the nature of decisions being made, potential harm to customers, the extent of human involvement in final decisions, transparency for affected customers and reliance on third-party data, models or systems.

64. Bermuda's market serves a large number of sophisticated institutional clients. Risk assessment frameworks should therefore consider impacts on other key stakeholders, including shareholders, business partners and employees, with appropriate consideration based on the nature of the business relationship and the specific sector context.

Implementation and Governance

65. Financial institutions should use these factors collectively to determine appropriate risk profiles and corresponding governance, oversight and control measures for their AI applications.

66. The BMA is also of the view that governance and risk management measures should scale proportionally across the risk tiers identified by financial institutions, with high-risk applications warranting comprehensive governance frameworks that include board-level oversight, regular independent validation, and enhanced documentation requirements. For smaller entities with less complex AI implementations, simplified frameworks that focus on key controls and alignment with the institution's broader governance framework may be appropriate.

67. The BMA further recognises that AI risk management should not exist in a silo but should be integrated with existing enterprise risk management frameworks. Financial institutions should incorporate AI risk considerations into their overall risk appetite statements, key risk indicators, and regular risk reporting to senior management and the board. This integration ensures AI risks receive appropriate attention as part of the institution's holistic risk governance structure, with clear accountability and oversight mechanisms that leverage existing control functions where appropriate while acknowledging the unique technical aspects of AI that may require specialised expertise.

68. This proportionality principle—scaling oversight intensity to match risk and institutional capacity—underpins every aspect of the BMA's AI governance framework. From inventory requirements to validation processes, each recommendation in this paper should be applied proportionately. This will ensure that

smaller and less complex financial institutions can participate in AI innovation without facing disproportionate compliance burdens while maintaining robust oversight where risks are material.

Discussion Questions for the Industry

To better understand the current landscape and future direction of AI adoption within Bermuda's financial services sector, please answer the following questions:

1. **Risk Assessment and Materiality Framework:** Do you agree with the proposed five-dimensional Risk Assessment Framework in paragraph 60? Given Bermuda's predominantly B2B institutional market context, are there additional risk dimensions or stakeholder considerations that should be incorporated to better capture the unique risk profiles of AI systems in your sector? Please provide specific suggestions for enhancement.

2. **Governance Integration and Board Accountability:** How does your organisation currently integrate AI risk management with existing enterprise risk management frameworks? Given the BMA's emphasis on ultimate Board accountability for AI outcomes, what specific challenges do you foresee in establishing appropriate Board-level AI literacy and oversight capabilities? What practical approaches have you found effective for ensuring clear accountability chains from technical implementation to Board governance?

3. **Proportionate Risk Management Implementation:** The BMA proposes that governance measures should scale proportionally across risk tiers, with simplified frameworks for smaller entities that have less complex implementations. What specific criteria or thresholds would you recommend for determining appropriate levels of governance intensity? How can the Three Lines of Defence Model be effectively adapted for AI systems while maintaining proportionality to risk and organisational complexity?

4. **Outcomes-Based Regulatory Approach:** The BMA's framework emphasises outcomes over prescriptive technical requirements. What specific outcomes or performance indicators do you believe would be most effective for demonstrating responsible AI governance in your sector? How can regulatory expectations be clearly defined while preserving the flexibility needed for innovation and sector-specific adaptations?

5. **Systemic Risk and Market-Wide Impacts:** Given Bermuda's integration into global financial markets, how do you assess the potential for your AI systems to contribute to market-wide or systemic risks? What mechanisms do you have in place to monitor for correlated behaviours with other market participants using similar AI systems, particularly during market stress conditions?

6. **International Regulatory Coordination:** How would international regulatory coordination and cooperation between Bermuda and other major financial centres enhance your ability to deploy AI systems globally? What specific areas of international regulatory coordination/cooperation would provide the most significant operational benefits for your cross-border activities?

AI Identification and Inventory

69. Without a comprehensive system to identify and take stock of AI systems and applications, financial institutions face significant governance blind spots where untracked AI systems operate without appropriate oversight or risk controls. This creates regulatory compliance risks, the potential for duplicated efforts, and the inability to assess aggregate AI exposure across the organisation. Undocumented AI systems may also introduce unmanaged dependencies and create operational vulnerabilities when systems fail or require

updates. Financial institutions should therefore identify and maintain comprehensive inventories of AI systems across all their business functions, documenting their purpose, technical specifications, risk classification, data sources and designated responsible parties to ensure appropriate oversight and governance. This inventory should include both internally developed and third-party AI solutions, tracking their lifecycle status, interdependencies with other systems, and changes in risk profile over time. Regular inventory updates enable financial institutions to maintain visibility over their AI ecosystem, facilitate risk assessment processes, support regulatory reporting requirements, and ensure that governance measures remain aligned with the evolving scope and complexity of AI deployments throughout the organisation. With comprehensive AI inventories in place, financial institutions can address the foundational elements underlying all AI systems –with robust data management processes.

Data Management

70. Poor data governance in AI systems can lead to biased outcomes, privacy violations and regulatory breaches that damage customer trust and institutional reputation. Inadequate data quality, lineage tracking or cross-border compliance frameworks expose financial institutions to legal liability and operational failures. Without robust data management, AI systems may make decisions based on incomplete, biased or outdated information, potentially resulting in unfair customer treatment and financial losses. Strong data governance throughout the AI lifecycle should address quality, operational relevance, appropriate bias minimisation, and compliance with privacy requirements, including Bermuda's Personal Information Protection Act 2016 (PIPA) and all applicable cross-border frameworks.

71. The cross-border nature of Bermuda's financial services necessitates having a clear view of PIPA with international privacy frameworks like the EU's General Data Protection Regulation, California Consumer Privacy Act/California Privacy Rights Act, and in all relevant jurisdictions where the financial institution operates. Financial institutions should implement a comprehensive data governance strategy for AI systems that maps cross-jurisdictional data flows, implements appropriate transfer mechanisms, maintains detailed processing records, and ensures consistent privacy standards regardless of data origin or processing location, while regularly reviewing AI privacy frameworks as global standards evolve.

Model Selection and Development

72. Inappropriate model selection or flawed development processes can result in AI systems that fail to meet business objectives or introduce unintended biases and errors. Without proper validation during development, financial institutions risk deploying systems that perform poorly in production environments or fail to comply with regulatory requirements. Poor development practices may also create models that lack necessary explainability or robustness for their intended use cases. AI models should be selected based on appropriateness for intended use cases, with financial institutions justifying their selection decisions and balancing complexity against explainability considerations. Model development processes should include appropriate testing for robustness, stability, and reliability across scenarios, maintaining comprehensive documentation covering goals, methodologies, and validation approaches.

73. Special attention should be paid to identifying potential errors, outliers and discrimination patterns. Development teams should also maintain separation from validation functions, with version control mechanisms to ensure auditability throughout the model lifecycle.

Human Oversight and Intervention

74. Insufficient human oversight creates accountability gaps where AI systems make critical decisions without appropriate human judgment or intervention capability. This can lead to automated decisions that harm customers, violate regulations or misalign with institutional values and risk appetites. Without clear escalation procedures and trained oversight personnel, financial institutions may be unable to detect or correct AI failures before they cause significant damage. Effective human oversight is essential for ensuring AI systems operate as intended and align with organisational values and regulatory requirements. Oversight does not always need to be real-time, but review cycles and escalation procedures must be swift enough to surface material issues before significant damage occurs. Financial institutions should establish clear guidelines for appropriate human intervention, particularly for high-risk AI applications that significantly impact customers or business-critical functions. Oversight personnel should receive specialised training on both technical aspects of AI systems (mainly focused on recognising anomalies in the model decision-making process) and ethical implications of AI-driven decisions.

75. Financial institutions should define explicit thresholds or scenarios that trigger mandatory human review of AI outputs. Once the thresholds and scenarios have been established, senior management should document the instances where human judgment overrides AI recommendations, including rationale and outcomes. Furthermore, the appropriate segregation of duties should be implemented between those developing AI systems and those providing oversight functions. Even well-trained staff may drift into complacency when oversight is seldom exercised, especially and paradoxically, in contexts with high automation performance. Financial institutions should build controls that force engagement—e.g., random challenge prompts, periodic shadow decisions, and Key Performance Indicator (KPI) tracking of human override rates to mitigate automation bias.

76. For systems with limited explainability, human oversight becomes particularly critical. It should be enhanced with non-AI/traditional controls and fallback procedures that ensure business continuity when AI systems fail to perform as expected.

77. Human oversight must extend beyond routine monitoring; it must encompass the authority and capability to halt an AI system when its behaviour threatens customers, market integrity or institutional solvency. The Authority, therefore, would expect all high-risk and critical AI applications to include an emergency override—colloquially, a *kill switch*—supported by clear trigger criteria, procedures and trained staff.

Automated Oversight and Model Verification Systems

78. While human-in-the-loop oversight remains essential, financial institutions may also implement "model-checking-models" approaches where appropriate. These systems employ independent AI models specifically designed to validate, monitor and detect anomalies in primary AI systems' outputs. Such configurations can provide continuous automated oversight while enhancing scalability.

79. Nonetheless, automated oversight systems can also create false confidence in AI governance or introduce additional complexity and failure points. Over-reliance on automated validation without human oversight could miss sophisticated errors or systematic biases that require human judgment to detect. These systems may also create cascading failures if the oversight models themselves become compromised or malfunction. Therefore, these systems should *complement*, rather than replace human oversight, particularly for high-risk applications where ultimate accountability still requires human judgment and intervention.

Explainability and Fairness

80. Lack of explainability in AI systems creates regulatory compliance risks and undermines customer trust, particularly when decisions significantly impact customer outcomes or access to financial services. Unfair AI systems can also perpetuate or amplify existing biases, resulting in discriminatory outcomes that violate anti-discrimination laws and harm an institution's reputation. Without proper fairness controls, financial institutions may inadvertently exclude specific customer segments or create unequal treatment, resulting in legal liability and regulatory sanctions. AI systems that handle customer-impacting decisions should maintain explainability to build trust and meet any relevant customer protection requirements, while demonstrating fairness to ensure equitable access to financial services. Financial institutions should develop explanation capabilities tailored to different stakeholders and establish robust processes to detect and mitigate unfair biases, with higher standards of explainability for systems that affect critical customer outcomes. Decisions should avoid being inaccurate, arbitrary or discriminatory, with special attention given to proxy variables that could inadvertently introduce bias. Financial institutions should document their fairness definitions, measurements and maintenance processes, ensuring alignment with any applicable laws or statutes and business objectives, while implementing appropriate customer notifications about AI usage throughout the financial service lifecycle.

81. The application of fairness principles should also reflect the nature of the business relationship, particularly in Bermuda's reinsurance and commercial insurance markets, where considerations may focus on contractual transparency, accuracy of risk assessments, and alignment with negotiated terms rather than common retail customer protection frameworks. The BMA also recognises that ethical considerations in such business transactions may extend beyond basic fairness considerations. These considerations should include particular legislation, corporate values alignment, transparency in decision-making processes, and cultural sensitivity, which takes into account Bermuda's position as a global hub serving diverse international clients. Financial institutions should therefore develop comprehensive AI ethics frameworks that articulate the institution's core principles and decision-making processes for resolving ethical dilemmas, with periodic review by both technical and non-technical stakeholders to ensure relevance as societal expectations and technology capabilities evolve.

Transparency and Disclosure

82. Insufficient AI transparency could create regulatory-compliance risks, weaken stakeholder trust and damage a financial institution's reputation. When customers remain unaware of how AI influences decisions such as investment advice, payment processing, or other services, financial institutions may become liable for opaque or unfair outcomes. A risk-based transparency approach may guide a financial institution's disclosure framework. For example, explanations should be matched to stakeholders' technical knowledge and legitimate interests. For higher-risk, customer-facing applications, meaningful explanations—supported by recognised explainability tools—could become essential. At the same time, excessive detail may expose proprietary algorithms, reduce competitive advantage or enable system manipulation. Financial institutions should seek an appropriate balance by protecting intellectual property and fraud-detection methods while still offering customers access to human oversight and providing supervisors with sufficient technical traceability. Stakeholders across the banking, insurance, investment, and payment sectors may require tailored disclosure practices.

83. A practical way to strike that balance is to adopt a three-layer transparency framework:

Layer	Purpose and illustrative content
Internal transparency	Financial institutions could develop model cards or factsheets that outline an AI system's purpose, data lineage, limitations and fairness indicators; financial institutions should consider maintaining an AI-incident register that is accessible to risk, audit and oversight functions.
Stakeholder-specific disclosure	Retail clients should receive plain-language notices, while institutional counterparties may request more technical annexes; regulators could be provided with fuller documentation such as code repositories, data-provenance records and validation summaries.
Public confidence — optional (large financial institutions only)	Systemically important or publicly listed firms may choose to include within existing public disclosures a concise statement of headline consumer-facing AI use-cases, key safeguards and customer-feedback mechanism channels. Smaller or non-public entities would not be expected to publish this layer.

Validation and Testing

84. Even well-designed AI systems can fail in unexpected ways. Without rigorous validation and testing, critical flaws may go undetected until they cause significant harm in production environments. Inadequate testing can miss performance issues, vulnerabilities or biases that only emerge in certain scenarios or with certain inputs. Systems might perform well on test data but fail catastrophically when exposed to real-world data distributions or edge cases not considered during development.

85. Independent validation provides an objective assessment of AI systems before deployment. The depth of validation should be proportionate to risk materiality as outlined in the previous section, with thorough documentation and remediation of identified issues. For third-party AI, alternative testing can verify performance in the institution's specific context. Validation techniques may include comparing model performance on unseen data that is available during development to performance on post-implementation data, to measure against expert review, or other methods appropriate to the technology. Testing should address interpretability, repeatability, robustness, reproducibility, traceability and model drift. Consequently, the testing documentation should also enable an independent party to reproduce the validation process and results. Organisations should ensure validation teams are independent from development teams to maintain objectivity.

Monitoring and Change Management

86. AI systems may degrade over time as real-world conditions change, requiring robust monitoring frameworks to detect performance degradation, data drift, concept drift, bias emergence and other anomalies unique to AI systems. Data drift occurs when the statistical properties of input data change (for example, customer demographics shifting after a merger), while concept drift happens when the relationship between inputs and outputs changes (such as fraud patterns evolving to bypass existing detection methods).

87. Financial institutions should establish clear thresholds for intervention, well-defined contingency plans, and structured change management processes for model updates with appropriate version control, documentation, and approval processes. For dynamic AI systems that update frequently, enhanced controls

are essential, including justification for automatic updating, clear definitions of permitted update parameters, and more stringent performance monitoring.

88. Financial institutions should also develop formal AI incident response procedures to address system failures, unintended consequences, or performance degradation, by defining what constitutes an AI incident and establishing clear escalation paths, key decision-makers and communication protocols, including with regulators and affected stakeholders. High-risk AI applications should undergo tabletop exercises to test procedure effectiveness before deployment, with regular reviews to incorporate emerging risk scenarios. Given the interconnected nature of financial services, financial institutions should consider how AI failures might propagate through the system and impact other market participants. All significant changes to production AI systems require review and approval by appropriate control functions, with comprehensive documentation including model code, data, hyperparameters and trained model weights to ensure complete traceability and accountability.

Third-Party Risk Management

89. Many financial institutions rely on third-party AI solutions, creating potential blind spots in risk management where the institution doesn't fully understand or control the underlying systems. Third-party AI could operate as a "black box" with unknown biases, vulnerabilities or limitations. Reliance on common vendors across the industry could create herding risk and concentration risk, where a single failure affects multiple financial institutions simultaneously. Contractual limitations might prevent proper testing or oversight.

90. For third-party AI, financial institutions should conduct due diligence proportionate to the system's materiality and risk profile. The intended use case, criticality to business operations, and potential impact on stakeholders should determine the depth of vendor assessment. For example, high-risk applications warrant comprehensive due diligence, contractual safeguards and ongoing monitoring, while lower-risk applications may require more streamlined oversight. Contractual terms for critical AI systems should provide appropriate audit rights and/or information access to ensure proper risk management.

Generative AI Considerations

91. Generative AI represents a new frontier with unique characteristics. These encompass the ability to create novel content and interact conversationally with customers, presenting novel risks not seen in traditional AI systems. Generative AI could produce hallucinations (confidently stated falsehoods), generate toxic or biased content, provide incorrect advice, or impersonate human representatives without proper disclosure. These systems can be particularly challenging to test comprehensively due to their extensive range of outputs.

92. To address these risks, financial institutions should implement mitigating controls, such as, but not limited to:

    - Human oversight for all critical decision-making points, input/output
    - Guardrails and filters that detect toxicity and biases
    - Grounding techniques such as Retrieval Augmented Generation (RAG), which combines AI with external knowledge sources to improve accuracy
    - Enhanced data security measures
    - Extensive validation testing

93. For Generative AI applications, financial institutions should focus on evaluating accuracy, relevance and bias using curated testing datasets specific to their requirements. Further, it is also advisable to implement more stringent validation requirements containing forward and edge case testing, to assess behaviour in production-like environments. Special attention should be paid to disclosure requirements when customers interact with generative AI systems.

Agentic AI Considerations

94. Agentic AI systems—capable of autonomous planning, decision-making, and action without continuous human guidance—represent an emerging frontier requiring specialised governance considerations. Unlike traditional AI systems, agentic AI can dynamically select tools, adapt workflows in real-time, and make consequential decisions independently. This autonomy introduces novel risks, including agency liability for AI-driven decisions, potential market instability from synchronised automated responses, tool choice hallucinations, and challenges in maintaining accountability chains. Financial institutions deploying agentic AI should implement enhanced oversight mechanisms like enhanced 'human-in-the-loop' governance models, real-time monitoring of autonomous decisions, and safeguards against cascading system-wide effects during market stress conditions.

Cybersecurity and Operational Resilience

95. AI systems introduce novel cybersecurity vulnerabilities that extend beyond traditional IT/cybersecurity frameworks. These include adversarial attacks designed to manipulate AI outputs, model poisoning during training phases, prompt injection attacks targeting generative AI, and unauthorised extraction of sensitive information through inference attacks, among others. As computing power and AI technology continue to evolve, the sophistication of adversarial attacks is also expected to evolve.

96. Financial institutions, therefore, should implement AI-specific security controls, including but not limited to:

    a) Regular adversarial testing and red team exercises specifically designed to probe AI vulnerabilities continuously

    b) Comprehensive monitoring for anomalies in model behaviour that might indicate compromise

    c) Automated alerts for significant deviations in AI performance or unexpected outputs

    d) Segmentation strategies that limit the impact of compromised AI components

    e) Enhanced access controls for training data, model parameters, and prompts

    f) Security assessment procedures that account for the unique attack vectors relevant to different AI architectures

97. Given Bermuda's prominence in complex risk transfer and cross-border transactions, financial institutions should conduct specialised threat modelling that considers both financially motivated actors and sophisticated nation-state adversaries that can target strategic financial infrastructure.

98. Additionally, operational resilience for AI systems requires robust planning that addresses unique failure modes not present in conventional software development. Financial institutions should develop AI-specific resilience measures that incorporate graceful degradation paths when AI components fail, redundant systems with different underlying architectures to prevent common-mode failures, and regular simulation of AI disruption scenarios.

99. Business impact assessments should also specifically evaluate the consequences of AI hallucinations, unexpected model drift, or sudden unavailability of third-party AI services. For critical financial functions, financial institutions should maintain alternative non-AI processes that can be activated when necessary, particularly for systems supporting catastrophe modelling, automated trading, or cross-border payment networks essential to Bermuda's financial services industry.

100. The BMA's Consultation Paper: *Operational Resilience Code and Guidance*[17] published in January 2025, provides a detailed discussion on the Authority's expectations from an outcomes perspective. This serves as a foundation for setting AI-specific guidelines that are helpful for financial institutions that are dedicated to building operationally resilient AI systems.

Insurance Coverage for AI-Related Risks

101. Even the most robust AI governance frameworks cannot eliminate virtually all risks, making insurance coverage a critical component of comprehensive AI risk management. Financial institutions should evaluate specialised insurance products designed to address AI-specific exposures that traditional cyber or professional liability policies may not cover adequately. These emerging insurance solutions can address risks, including algorithmic bias claims, model failure, AI hallucinations that cause financial loss, intellectual property infringement, and third-party AI vendor liability. When evaluating coverage, financial institutions should consider policy features including coverage trigger definitions, exclusions related to intentional discrimination, contractual liability limitations, and the scope of coverage for regulatory investigations.

102. Financial institutions should work closely with insurance partners who understand both AI systems and the unique risk profiles of Bermuda's financial services ecosystem. Given the rapidly evolving nature of AI risks, financial institutions should regularly reassess the adequacy of their coverage as their AI deployments mature. However, financial institutions should not view insurance as a substitute for strong governance, but as a complementary control that provides financial protection to recover from governance failures. This approach aligns with the BMA's commitment to operational resilience and ensures that financial institutions can maintain financial stability even when faced with significant AI-related incidents or disruptions.

Discussion Questions for the Industry

To better understand the current landscape and future direction of AI adoption within Bermuda's financial services sector, please answer the following questions:

1. **Sector-Specific Adaptations and Cross-Sectoral Applications:** How should the proposed framework be adapted to address the unique characteristics of your specific sector within Bermuda's financial services ecosystem? What emerging AI risks or use cases specific to your sector require additional consideration beyond the framework outlined in this paper?

2. **Enhanced Third-Party AI Risk Management:** What specific challenges has your organisation encountered in managing third-party AI solutions, particularly regarding: (a) due diligence and ongoing oversight of AI service providers; (b) contractual arrangements that ensure adequate regulatory access and control; and (c) managing concentration risk where multiple financial institutions rely on the same AI providers? How do you address the complexity of third-party relationships and subcontractor dependencies in your AI supply chain?

[17] BMA (2025), *Consultation Paper: Operational Resilience and Outsourcing Code.*

3. **Stakeholder-Specific Transparency and Communication:** How does your institution currently approach transparency and disclosure regarding AI usage in customer-facing decisions? Please describe your current practices for: (a) communicating with sophisticated institutional clients about AI usage; (b) providing appropriate disclosures to retail customers; (c) maintaining documentation for regulatory supervisors; and (d) balancing transparency obligations with protecting proprietary systems and maintaining operational security.

4. **Cybersecurity and Operational Resilience:** What AI-specific cybersecurity vulnerabilities and operational resilience challenges have you identified that may not be adequately addressed by traditional IT security frameworks, as well as the current operational cyber risk management codes of practice applicable to insurance, digital asset business, corporate service providers, trust companies, money service businesses, investment businesses and fund administration providers, banks and deposit companies and the proposed Operational Resilience Code of Conduct and Guidance Note? How do you approach threat modelling for AI systems, particularly considering both financially motivated actors and sophisticated nation-state adversaries targeting Bermuda's strategic financial infrastructure? What contingency measures do you have in place for AI system failures or security compromises?

5. **Agentic AI Deployment and Governance:** Does your organisation currently deploy or plan to deploy agentic AI systems (AI capable of autonomous decision-making, planning, and tool selection without continuous human oversight)? If so, what specific governance challenges have you encountered in maintaining accountability for autonomous AI decisions? What regulatory guidance would be most valuable for managing risks from agentic AI systems, particularly regarding agency liability, market stability impacts and the oversight of real-time autonomous adaptations?

6. **Implementation Practicalities and Resource Allocation:** What specific implementation challenges do you anticipate in your organisation, particularly regarding resource allocation, timeline management and integration with existing systems? How do these challenges differ between internal AI deployments versus third-party AI solutions?

Overall Considerations

103. While the framework proposed and outlined in the previous section provides a comprehensive foundation for AI governance across financial services, different sectors within Bermuda's diverse financial ecosystem face unique challenges and opportunities for consideration.

104. The BMA acknowledges that implementing robust AI governance frameworks involves costs and resource commitments that may present challenges, particularly for smaller and less complex financial institutions and those new to AI adoption. To address this, the proportionality principle extends not only to the scope of controls, but also to implementation timelines, resource allocation, and the flexibility to use collaborative approaches. The Authority also recognises that effective regulation must be practical—achievable by the diverse range of financial institutions it supervises.

105. Building on these overall considerations regarding costs and sector-specific challenges, the following section provides specific guidance on practical implementation approaches that enable financial institutions to achieve effective governance outcomes while managing resource constraints.

106. Effective implementation of this framework requires the careful consideration of both proportionality in application and practical execution strategies that enable all financial institutions—regardless of size or

complexity—to achieve appropriate governance outcomes while fostering continued innovation in Bermuda's dynamic financial services marketplace.

Implementation Considerations

107. While core AI governance principles apply across all financial sectors, implementation should reflect the proportionality principle, considering institutional scale, complexity, AI maturity and customer relationship types. The BMA recognises that implementing comprehensive AI governance frameworks involves significant resource commitments and expertise requirements that may challenge smaller financial institutions or those in the early stages of AI adoption. This section provides practical guidance on achieving effective governance outcomes through proportionate and phased approaches that balance innovation with appropriate risk management.

Proportionality Principle

108. Consistent with the BMA's general approaches to supervising financial sectors, the proportionality principle underpins AI governance expectations across all financial institutions, recognising that organisations vary significantly in nature, scale, complexity, AI maturity and risk appetite. The BMA proposes a tiered approach where oversight intensity corresponds to an AI system's potential impact and materiality, with streamlined requirements for lower-risk applications and enhanced controls for high-risk systems.

109. In applying proportionality, financial institutions should explicitly consider the nature of customer relationships (retail versus institutional), with Bermuda's significant institutional market context allowing for different approaches to transparency, explainability and customer protection compared to retail-focused jurisdictions. Financial institutions should conduct comprehensive risk-based assessments that evaluate both the complexity of AI technology and its business criticality to determine appropriate controls, enabling smaller entities to innovate without disproportionate compliance burdens while ensuring robust oversight for applications that could significantly impact financial stability, market integrity or customer outcomes.

Practical Implementation Strategies

110. Financial institutions should consider phased implementation approaches that prioritise governance efforts based on risk materiality and organisational readiness, with the initial phases focusing on establishing foundational governance structures, conducting comprehensive AI inventories, and implementing controls for the highest-risk applications before expanding to lower-risk systems. This staged approach allows financial institutions to build expertise and refine processes while managing implementation costs and resource constraints.

111. Financial institutions may benefit from collaborative approaches such as shared industry utilities for common AI governance functions, joint procurement of specialised expertise, or participation in industry working groups to develop sector-specific best practices. Financial institutions should leverage existing enterprise risk management frameworks where possible. They should adapt established processes rather than create entirely separate AI governance structures. However, financial institutions must recognise that AI's unique technical aspects may require specialised expertise or enhanced procedures.

Resource Capability Development

112. Effective AI governance requires financial institutions to develop appropriate technical and risk management capabilities across multiple organisational levels. Board members and senior management need sufficient AI literacy to provide meaningful oversight, while operational staff require training on the AI-specific risks and controls relevant to their roles. Financial institutions should establish systematic training programmes that evolve with technological developments and organisational AI maturity.

113. Where internal expertise is limited, financial institutions may engage external specialists for specific functions such as independent validation, specialised auditing or technical advisory services, provided such arrangements include appropriate oversight mechanisms to ensure external providers understand the institution's risk appetite and regulatory obligations. Documentation and knowledge transfer requirements should ensure that critical insights are retained within the organisation while maintaining ultimate accountability for AI outcomes.

Monitoring and Continuous Improvement

114. AI governance frameworks should be viewed as dynamic systems that require regular review and adaptation as technology evolves, organisational capabilities mature, and risk profiles change. Financial institutions should establish periodic assessments of their governance effectiveness, incorporating lessons learned from AI incidents, regulatory developments and industry best practices to ensure governance measures remain proportionate to evolving risks and business objectives.

115. The BMA expects financial institutions to maintain comprehensive documentation of their AI governance decisions, implementation approaches and ongoing assessments to support supervisory review and demonstrate compliance with outcomes-based expectations. This documentation should enable both internal stakeholders and external reviewers to understand how governance measures align with the institution's risk profile and regulatory obligations.

Discussion Questions for the Industry

To better understand the current landscape and future direction of AI adoption within Bermuda's financial services sector, please answer the following questions:

1. **Proportionality Criteria and Thresholds:** How should the BMA establish clear thresholds or criteria to determine proportionate regulatory requirements for different AI applications, considering factors such as the sector you currently operate in, materiality, autonomy level, the type of customers and perceived market impact, as well as overall market stability?

2. **Regulatory Documentation and Attestation Requirements:** What regulatory documentation or attestation requirements would be appropriate for different categories of AI risk, to balance the need for effective oversight with the goal of encouraging innovation across Bermuda's diverse financial ecosystem?

3. **Market Outcomes and Success Metrics:** Looking beyond compliance, what specific market outcomes would indicate successful implementation of the BMA's AI governance framework? What key performance indicators would you recommend for evaluating the framework's effectiveness?

4. **Regulatory Innovation and Future-Readiness:** How can the BMA best support innovation in AI applications that serve sophisticated institutional clients while maintaining appropriate safeguards? What emerging AI systems or applications should the BMA be preparing for in developing its framework?

IX. Conclusion

116. The rapid evolution of AI represents both a transformative opportunity and a regulatory imperative for Bermuda's financial services sector. This DP proposes a principles-based, outcomes-focused regulatory approach that supports technological innovation while maintaining the robust regulatory standards that have established Bermuda's reputation as a premier financial jurisdiction. Rather than prescriptive technical requirements, the BMA's framework emphasises board accountability, proportionate risk management, and integration with existing regulatory structures to create a coherent supervisory ecosystem.

117. The proposed approach also recognises Bermuda's unique market characteristics—sophisticated institutional clients, complex cross-border arrangements, and diverse financial sectors—while maintaining consistent core principles across all applications. By focusing on governance outcomes rather than implementation details, financial institutions retain flexibility to develop AI systems tailored to their specific business contexts while ensuring appropriate safeguards. This framework builds upon existing regulatory foundations rather than creating separate AI-specific obligations, reducing compliance complexity while addressing the novel challenges these technologies present.

118. The BMA acknowledges that effective AI governance extends beyond individual institutional risks to encompass market-wide considerations, including concentration and herding risks from common service providers, the potential for correlated behaviours, and possible cascading failures across interconnected systems. The Authority's supervisory approach will monitor these systemic dimensions through ongoing market engagement, information sharing with international counterparts, and coordination with major technology providers to ensure Bermuda's financial ecosystem remains resilient as AI adoption accelerates.

119. The BMA is committed to positioning Bermuda as a jurisdiction where financial institutions can confidently leverage AI systems to enhance efficiency, improve client outcomes, and develop innovative products and services, while maintaining the high standards of stability, integrity and customer protection. While international consistency is important, particularly for cross-border operations, Bermuda has an opportunity to develop a distinctively balanced approach that reflects our market position and principles-based regulatory philosophy. The rapidly evolving nature of AI systems necessitates close collaboration between the BMA, regulated entities and other stakeholders to develop effective oversight mechanisms that remain adaptable to technological changes.

Next Steps

120. Following the consultation period, the BMA will analyse stakeholder feedback, develop specific proposals through targeted consultation papers, conduct focused industry workshops, publish an implementation roadmap with clear timelines, and develop supporting resources for the market, including practical guidance, case studies and capability-building initiatives to enhance AI literacy across the financial services sector.

121. This paper represents the beginning of a dialogue rather than its conclusion, and the BMA invites all stakeholders to submit detailed responses by 30 September 2025 to help shape a framework that balances technological innovation with the appropriate safeguards.

Call to Action

122. The Authority particularly encourages feedback on:

a) The appropriateness of the proposed risk management framework for your specific sector;

b) Practical implementation challenges you foresee with the proposed approach;

c) Additional risk considerations unique to the Bermuda market;

d) Alternative regulatory approaches that might achieve the same objectives;

e) Specific areas where a pragmatic regulatory approach would encourage beneficial innovation; and

f) Whether the proportionality principle adequately addresses the needs of different-sized institutions.

123. The industry's insights will directly inform the development of a more tangible proposal for the development of the AI Policy in Bermuda, helping to maintain its position as a forward-thinking jurisdiction that balances technological innovation with appropriate safeguards.

124. Please submit responses using the online survey form or directly to policy@bma.bm with the subject line Discussion Paper – *The Responsible Use of Artificial Intelligence in Bermuda's Financial Services Sector*.

X. Glossary of Terms

Adversarial Attacks: Deliberate attempts to manipulate AI systems by introducing carefully crafted inputs designed to cause incorrect outputs or decisions.

Agentic AI: Advanced artificial intelligence systems capable of autonomous planning, decision-making and action execution without continuous human guidance. Unlike traditional AI, which responds to specific prompts, agentic AI can independently perceive environments, reason about complex scenarios, select appropriate tools, and adapt strategies in real-time to achieve specified objectives.

Agency Liability: Legal responsibility that arises when AI systems are empowered to make decisions or take actions traditionally performed by human employees. Under agency liability principles, both the deploying organisation and potentially the AI provider may be held directly accountable for autonomous AI decisions and their consequences.

AI (Artificial Intelligence): Technologies that perform functions typically requiring human intelligence, including learning, reasoning, problem-solving and decision-making.

Algorithm: A set of rules or instructions given to an AI system to help it learn from data and make decisions.

Autonomy: In AI systems, the ability to operate and make decisions with varying degrees of human intervention.

Bias: Systematic errors or unfairness in AI systems that can lead to discriminatory outcomes against certain groups or individuals.

BMA: The Bermuda Monetary Authority, also known as the Authority or BMA, is the integrated regulator of Bermuda's financial services sector.

Concept Drift: Changes in the underlying relationship between input variables and target outcomes over time, requiring AI models to be retrained or adjusted.

DABA: The Digital Asset Business Act 2018, Bermuda's regulatory framework for digital asset businesses.

Data Drift: Changes in the statistical properties or distribution of input data over time, potentially making AI models less accurate or reliable than when originally trained.

EU AI Act: The European Union's comprehensive regulatory framework for AI that classifies AI systems based on risk levels.

Explainability: The ability to explain how and why an AI system made a particular decision in terms that are understandable and appropriate to different stakeholders, including customers, regulators and technical staff.

FATF: Financial Action Task Force, an intergovernmental organisation that develops policies to combat money laundering and terrorism financing.

Generative AI: AI systems capable of creating new content (i.e., text, images, audio) that was not explicitly programmed, including Large Language Models (LLMs).

Governance: Comprehensive frameworks, policies, processes and oversight mechanisms for ensuring AI systems operate as intended, align with organisational values and regulatory requirements, and maintain appropriate accountability structures.

Hallucinations: In generative AI, the production of confident but factually incorrect or nonsensical outputs that appear plausible.

Herding Risk: The systemic risk that arises when multiple financial institutions rely on similar AI systems, vendors or methodologies, leading to correlated decision-making and synchronised market behaviours that can amplify market movements or cause cascading failures across the financial system during stress conditions.

Human-in-the-Loop: AI system design that incorporates meaningful human oversight and intervention capabilities at critical decision points.

IAIS: International Association of Insurance Supervisors, a global standard-setting body for insurance supervision.

IOSCO: International Organization of Securities Commissions, a global standard-setting body for securities markets.

KYC/AML: Know Your Customer/Anti-Money Laundering, regulatory processes designed to verify customer identity and prevent illegal activities.

LLM (Large Language Model): A type of advanced AI model trained on vast amounts of text data that can understand, generate and manipulate human language. LLMs underpin many generative AI applications and can perform tasks such as drafting documents, answering questions, summarising content and translating languages with human-like capabilities.

Machine Learning (ML): A subset of AI where systems improve their performance on tasks through experience without being explicitly programmed.

Model Drift: The gradual deterioration of an AI model's performance over time as the relationship between input variables and target outcomes changes, often due to shifts in real-world conditions that were not represented in the original training data. There are two main types: ‘data drift’ (changes in input distribution) and ‘concept drift’ (changes in the underlying relationships).

Model Poisoning: A cyberattack where malicious data is introduced during the AI training process to compromise the model's integrity and performance.

Model Validation: The process of evaluating an AI model's performance, stability and appropriateness for its intended use.

NAIC: National Association of Insurance Commissioners, the US standard-setting organisation for insurance regulation.

NIST: National Institute of Standards and Technology, a US agency that develops technology standards, including the AI Risk Management Framework.

PIPA: Personal Information Protection Act 2016, Bermuda's legislation governing data privacy.

Principles-Based Regulation: A regulatory approach that focuses on broad principles rather than prescriptive rules, allowing flexibility in implementation.

Prompt Injection: A type of cyberattack targeting generative AI systems where malicious instructions are embedded in user inputs to manipulate system behaviour.

Proportionality: The principle that regulatory requirements should be appropriate to an entity's size, complexity and risk profile.

Proxy Variables: Data elements that indirectly represent protected characteristics and could lead to discriminatory outcomes if not properly managed.

RegTech (Regulatory Technology): Technology solutions designed to help financial institutions meet regulatory requirements efficiently and effectively. RegTech applications leverage technologies such as AI, machine learning and data analytics to streamline compliance processes, enhance regulatory reporting and improve risk management.

Retrieval Augmented Generation (RAG): A technique that enhances generative AI by combining it with external knowledge sources to improve accuracy and reduce hallucinations.

Risk-Based Approach: A regulatory and management philosophy that focuses resources and requirements on areas presenting the highest risks, with control measures proportionate to the potential impact and likelihood of adverse outcomes.

SupTech (Supervisory Technology): The use of innovative technology by supervisory agencies to support supervision.

Third-Party Risk: Risks arising from the use of external vendors or providers of AI systems or components.

Three Lines of Defence Model: Risk management framework dividing responsibilities between business units (first line), risk/compliance functions (second line) and internal audit (third line).


BMA
BERMUDA MONETARY AUTHORITY