BERMUDA MONETARY AUTHORITY

Friday, 14 August, 2026

CONSULTATION PAPER

The Responsible Use of Artificial Intelligence in Financial Services in Bermuda

Comments to be received by 30 October 2026

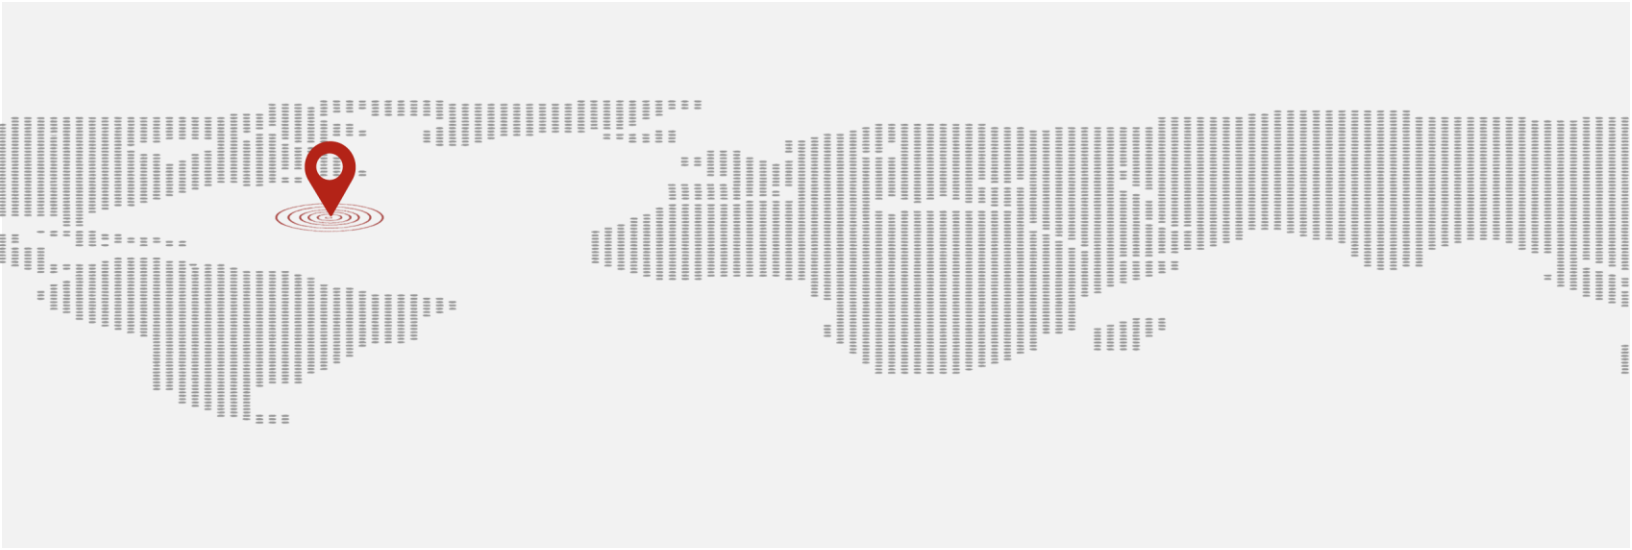

Table of Contents

PART A - CONSULTATION PAPER

I. INTRODUCTION

1. The global financial services landscape is undergoing a profound transformation driven by advances in Artificial Intelligence (AI). AI adoption is being driven by improvements in computing capacity, cloud infrastructure, data availability, foundation models and user-friendly interfaces, alongside competitive pressure, cost reduction, demand for faster service and the increasing complexity of financial crime, cyber threats and regulatory obligations. Current and emerging AI applications within financial services range from established machine-learning models used in underwriting, credit assessment, fraud detection and market surveillance to generative AI (GenAI) tools used for document analysis, coding and customer engagement, and emerging agentic AI systems capable of planning and executing multi-step tasks. Properly governed, these applications can improve efficiency, risk identification, customer and client service, compliance and operational resilience. They can also introduce or amplify model, data, conduct, cyber, outsourcing, operational and financial stability risks.
2. Regulated entities are also increasingly exposed to AI through enterprise productivity tools and embedded third-party software, which may lower barriers to adoption but can create shadow AI, data leakage, uncontrolled model changes and unclear accountability where governance does not keep pace. Responsible adoption therefore requires regulated entities to manage both the risks of using AI and, where relevant, the strategic or operational risks of failing to modernise capabilities as market practices, threats and supervisory expectations evolve.
3. As AI adoption across Bermuda's regulated financial services sector continues to evolve, the Bermuda Monetary Authority (Authority or BMA) considers it appropriate to provide guidance on the responsible use of AI by entities it regulates and supervises.
4. This Consultation Paper (CP) builds on the Authority's July 2025 Discussion Paper (DP) on the Responsible Use of AI in Bermuda Financial Services Sector. It is informed by the review and analysis of stakeholder feedback, cross-sector supervisory analysis, international benchmarking and the Authority's continued engagement with global regulatory developments.
5. The CP has four principal objectives:
    a. to illustrate some of the potential opportunities and risks associated with the use of AI by regulated entities;
    b. to summarise the Authority's consideration of stakeholder feedback on the DP;
    c. to consult on a proposed Guidance Note; and
    d. to maintain an open, evidence-based dialogue with the market as AI use expands.
6. Responses should be submitted by 30 October 2026 to policy@bma.bm

II. EXECUTIVE SUMMARY

7. Stakeholder feedback on the Authority's DP broadly supported the Authority's proposed principles-based, outcomes-focused and proportionate approach. Respondents suggested that AI-related risks are substantially addressed by Bermuda's existing regulatory frameworks and indicated a preference against the introduction of a separate AI approval regime. At the same time, stakeholders

requested clarity regarding the extent to which existing regulatory requirements apply to the use of AI.

8. Following its consideration of stakeholder feedback, the Authority published a Stakeholder Letter on 9 February 2026, conducted a cross-sector mapping of its supervisory framework and reviewed relevant international laws, standards and supervisory publications. This review identified common themes across emerging regulatory and supervisory approaches to AI, including governance and accountability, the integration of AI-related risks into enterprise risk management frameworks, proportionality, model lifecycle oversight, data governance (including the ethical sourcing of data), validation and testing, human oversight, transparency, monitoring, cyber resilience and the management of risks associated with third-party providers.
9. The Authority's assessment indicated that many of these themes were reflected, either directly or indirectly, within Bermuda's existing regulatory and supervisory frameworks. However, the evolving use of AI raised questions regarding the interpretation and application of such existing obligations. Consistent with stakeholder feedback, the Authority considers that the provision of regulatory clarity through supervisory guidance is an appropriate approach at this stage. The Authority will continue to monitor market developments, international standards and supervisory experience, and will assess whether targeted amendments to regulatory requirements may be warranted in the future.
10. The Authority's approach is also calibrated to Bermuda's market structure. Bermuda's financial services sector includes substantial international, wholesale and institutional business, including insurance and reinsurance, captives, insurance-linked securities, investment funds, digital asset business, trust, corporate service provider and fund administration activities. Many BMA-regulated entities operate through cross-border groups, global service providers, shared platforms and specialist third-party vendors. The Authority's approach therefore focuses on the regulated activity, use-case materiality, affected stakeholders and existing legal obligations, rather than applying a one-size-fits-all AI taxonomy or importing retail-focused requirements without appropriate adaptation.
11. The proposed Guidance Note (GN) treats the AI use case, rather than the technology label or the size of the entity alone, as the primary unit of analysis. The nature and intensity of governance, risk management documentation, oversight and assurance should reflect both the materiality and risk of the use case, together with the nature, scale, complexity, business model and resources of the regulated entity. Proportionality does not require smaller or less complex entities to replicate the governance structures or specialist functions of larger institutions where equivalent outcomes can be achieved through simpler arrangements.
12. For the purposes of the GN, materiality includes the importance of the AI use case to the entity's viability, critical operations, important business services and ability to meet its legal and regulatory obligations. Risk, on the other hand, includes the potential for harm to a regulated entity's customers, investors, counterparties, markets or other affected stakeholders. A lower-materiality use case may still present a higher risk where, for example, it involves sensitive information.
13. The proposed GN is organised around five outcomes: (i) **accountable governance and risk ownership**; (ii) **proportionate lifecycle management**; (iii) **reliable, explainable and appropriately overseen AI outcomes**; (iv) **secure and resilient deployment**, including third-party

arrangements; and (v) **adequate evidence to enable internal assurance and supervisory review**. Sector-specific application and considerations will remain sensitive to Bermuda's predominantly international, wholesale and institutional market while preserving appropriate protections for retail customers, investors, policyholders and other affected parties.

14. The Authority seeks feedback on the proposed GN – in particular, whether it provides sufficient clarity as to the application of existing regulatory frameworks, whether the risk-based proportionality framework is operationally feasible for different sectors of Bermuda's market and whether any important sector-specific or emerging risks require further consideration before finalisation.

III. SCOPE AND LEGISLATIVE BASIS

Scope of Consultation

15. Part B of this CP contains the text of the draft proposed GN on the Responsible Use of AI in Bermuda's Financial Services Sector. The GN seeks to clarify how existing legislative, regulatory and supervisory requirements may apply where a regulated entity develops, procures, deploys or otherwise uses AI as part of its business and operating model where such use may affect its regulated activities, compliance with applicable laws, customers, policyholders, clients, investors, counterparties, beneficiaries, digital asset acquirers, market participants, critical operations or the safety and soundness of the entity.
16. The GN is intended for use across the financial services entities regulated and supervised by the Authority under legislation administered by the Authority ('regulated entities'). It applies where an AI system is:
    a. developed internally;
    b. procured from or operated by a third-party service provider;
    c. provided by, shared with or operated by another member of the regulated entity's group;
    d. embedded within another product, service, platform, model, application or business process;
    e. hosted on the regulated entity's own infrastructure or on external infrastructure; or
    f. used directly by the regulated entity or by a person acting on its behalf.
17. Regulated entities using AI remain responsible for complying with their applicable legal and regulatory obligations. The use of a third-party, group-provided or externally hosted AI system does not transfer or diminish a regulated entity's accountability for the resulting regulated activity, decision, process or outcome.

Legislative Basis and Status

18. The Authority proposes this GN in the exercise of its functions under the Bermuda Monetary Authority Act 1969 and the sectoral legislation it administers. Regulated entities remain subject to their applicable sector-specific statutory and regulatory requirements.
19. For the purposes of the CP and the proposed GN, references to the applicable regulatory framework include the relevant Acts, regulations, rules, codes of conduct, codes of practice, prudential

standards, guidance notes, statements of principles, licence conditions, directions and other supervisory requirements applicable to the relevant regulated entity or regulated activity.

20. The GN does not create new statutory or regulatory obligations, prescribe a single method of compliance or establish a separate approval regime for the use of AI. It describes how existing obligations and supervisory expectations may apply where an entity uses AI and identifies illustrative practices that an entity may consider when managing material AI-related risks.
21. Where an applicable Act, regulation, rule, code of conduct or code of practice establishes a specific requirement, that requirement prevails. The GN should be read alongside, and does not amend or replace, the Authority's existing frameworks relating to governance, conduct, cyber risk, operational resilience, outsourcing, AML/ATF and sanctions compliance.
22. The appropriate approach will depend on the nature, scale, complexity, autonomy and potential impact of the AI use case. An entity may adopt alternative controls where these achieve outcomes consistent with its applicable legal and regulatory obligations.

IV. THE BERMUDA CONTEXT

23. Bermuda's approach to the supervision of AI-related risks should be understood in the context of its role as an international financial centre. The market includes a significant concentration of wholesale, institutional, cross-border and group-connected business. AI use in this context may therefore affect a range of stakeholders and regulated entities' own safety and soundness in different ways. Accordingly, the Authority considers that AI governance should be calibrated to the use case, regulated activity and relevant stakeholders.
24. Given this wholesale context and the reliance on international group-provided systems, expectations around transparency and fairness should not be imported without adaptation from retail-focused regimes. Accountability remains essential, and because AI may influence high-value activities with significant prudential or market consequences – such as underwriting, catastrophe modelling, reserving, portfolio management, risk transfer and AML/ATF monitoring – institutional counterparties, boards, auditors and supervisors must retain sufficient information to independently challenge those outcomes.
25. Retail and customer-facing use and impact cases will warrant attention to effective disclosure, explainability, contestability, redress and fair outcomes. A regulated entity should recognise and take into account the status of a relevant affected party and the duties owed when considering how to assess and apply regulatory requirements in an AI context.
26. BMA-regulated entities frequently operate through cross-border and group arrangements. AI systems may be developed, hosted, updated or supported outside of Bermuda, process data across multiple jurisdictions and rely on common foundation models or infrastructure. Boards and management should retain oversight, provide challenge and remain ultimately accountable for the outcomes.
27. The concentration of cloud infrastructure, specialised hardware, foundation models and AI platforms may also create common dependencies across the market and regulated entities should consider material dependencies on common AI providers, models, infrastructure or group platforms within existing outsourcing, operational resilience, concentration risk and business continuity arrangements.

V. APPLICATIONS AND OPPORTUNITIES

28. AI use across Bermuda's financial services sectors offers substantial opportunities to enhance efficiency, risk management, compliance, customer and client service, and operational resilience. The following examples are illustrative and do not imply that AI is appropriate for every use case:

Sector or function	Illustrative applications	Potential benefits
Insurance and reinsurance	Underwriting support, catastrophe and exposure modelling, pricing, claims triage, fraud detection, portfolio optimisation, document extraction	Faster analysis, improved consistency, enhanced risk segmentation, reduced manual processing
Banking and credit	AML/KYC, fraud monitoring, credit decision support; customer service, treasury and liquidity analysis	Improved detection, faster onboarding, enhanced decision making, greater operational efficiency, improved customer outcomes
Investment business and funds	Research and summarisation, portfolio analytics, trading surveillance, compliance review, investor reporting	Processing of large datasets, improved surveillance, faster reporting and analysis
Digital assets and payments	Blockchain analytics, transaction monitoring, cybersecurity, payment routing, smart-contract review	Real-time monitoring, fraud prevention, enhanced compliance and efficiency
Trust, corporate and administrative services	Document review, beneficial ownership analysis, due diligence, regulatory and tax reporting support	Reduced processing time, improved consistency, enhanced monitoring
Cross-sector risk and operations	Coding assistance, knowledge retrieval; cyber defence, scenario generation; incident analysis, regulatory change monitoring	Productivity, improved risk identification, stronger resilience and compliance

29. To ensure that the benefits of AI are realised safely, the Authority's supervisory focus extends beyond the mathematical model to its operational environment. A high-performing model can still produce inappropriate outcomes if it is trained on poor-quality data, used outside its intended context, inadequately monitored or integrated into a process without meaningful human challenge. Accordingly, robust governance requires attention to the entire use-case lifecycle, rather than the technical model alone.

VI. KEY RISKS AND CHALLENGES

30. In many cases, AI impacts the speed, scale, complexity or opacity of risks. The Authority considers the following risk areas particularly relevant to the use of AI:

    a. governance and accountability;

    b. model and performance risks;

c. data quality, privacy and bias risks;
d. explainability and transparency;
e. risks arising from ineffective human oversight or excessive automation;
f. consumer protection and conduct risks;
g. cyber and technology risk;
h. operational resilience and third-party risk;
i. market integrity; and
j. financial stability, concentration and systemic risks.

31. AI-specific features can make traditional controls less effective. Adaptive systems may change after deployment; GenAI can produce plausible but inaccurate outputs; agentic AI can take a series of autonomous actions; complex models may be difficult to explain; and common vendors may change underlying models without adequate notice. These features may require proportionate enhancements to existing frameworks.
32. 'Shadow AI', including the unauthorised or inappropriate use of AI by employees, is also a growing governance challenge. Controls should address tool approval, access, data classification, monitoring, training and escalation.
33. AI systems provided by third-party vendors may present challenges for smaller entities and Bermuda subsidiaries that may have limited negotiating leverage. It may not be viable for regulated entities to obtain proprietary source code, model weights or training data. Examples of expected outcomes relate to sufficient assurance, information, testing, contractual protection, monitoring and compensating controls to manage the use case within risk appetite.

VII. BMA WORK AND STAKEHOLDER FEEDBACK

BMA Survey and Internal Application

34. The Authority's 2022 insurance sector AI and machine-learning survey found that AI adoption was already occurring across the market, with prominent use cases in cybersecurity, underwriting and claims. It also identified challenges relating to written strategies, expertise, explainability and auditability. The subsequent mainstream adoption of GenAI and the emergence of agentic AI have increased the urgency of those issues.
35. The BMA has also continued to develop its own data science, analytics and supervisory technology capabilities. This experience informs the Authority's approach to practical, risk-based oversight and reinforces the importance of sound governance, data management, validation and human accountability when AI is used in supervisory or regulated contexts.

Stakeholder Feedback on the DP

36. Stakeholders broadly supported an outcomes-based, principles-led and technology-neutral approach. There was strong support for integrating AI into existing enterprise risk management and governance frameworks rather than creating AI frameworks that may prove duplicative.

Respondents also emphasised proportionality, international alignment and the need to recognise Bermuda's institutional and wholesale market.

37. Stakeholders requested greater clarity as to the scope and definition of AI, materiality thresholds, board AI literacy, use-case inventories, human oversight, explainability, validation, third-party arrangements, GenAI controls and the supervisory evidence that regulated entities should be able to provide.

38. Stakeholders also cautioned that rigid requirements could be counterproductive. Examples included requiring human review of every low-risk output, imposing detailed inventory requirements on basic productivity tools, demanding technical access that vendors will not provide, or applying retail disclosure and fairness concepts without regard to wholesale contractual relationships.

International Benchmarking and Policy Analysis

39. As part of its analysis, the Authority benchmarked its proposed approach outlined in the DP against an updated list of relevant international laws, standards, supervisory statements and industry guidance. A list of international sources and abbreviations is included in Annex E.

40. The benchmarking demonstrated strong international convergence in the following areas:

    a. ultimate accountability remains with the regulated entity, its board and senior management;

    b. AI controls should be integrated into governance, enterprise risk management and the three lines of defence;

    c. controls should be proportionate to use-case materiality and risk;

    d. AI use cases should be identified, documented and governed throughout their lifecycle;

    e. data quality, provenance, security and lawful use and ethical sourcing are foundational;

    f. material AI should be validated, tested and monitored for performance degradation and drift;

    g. human oversight should be meaningful and appropriate to autonomy, complexity and impact;

    h. explainability and transparency should be proportionate and tailored to stakeholders;

    i. third-party reliance does not transfer accountability and requires appropriate assurance;

    j. AI-related cyber, operational resilience, concentration and systemic risks should be monitored; and

    k. fairness, bias mitigation and the prevention of unintended discriminatory outcomes should be considered throughout the AI lifecycle.

41. While supervisory and regulatory approaches differ across jurisdictions, the benchmarking identified a number of areas where there is strong international convergence. These themes have informed the development of the BMA's proposed principles-based and proportionate approach to AI governance and risk management. While the BMA is not proposing to replicate any particular international regime, the benchmarking exercise has been used to identify practices and expectations that appear consistently across leading standards, guidance and supervisory frameworks and that may be relevant to Bermuda's financial services sector.

VIII. PROPOSED REGULATORY APPROACH

42. The proposed GN is relevant for entities across BMA-regulated sectors. It should be interpreted in a manner consistent with the scope and purpose of the underlying requirements within each sector.

43. The proposed approach is based on the following principles:

    a) **technology neutrality** - regulatory expectations should be driven by risk and outcomes rather than the technology itself. Equivalent risks should be treated consistently regardless of the technology used;

    b) **accountability** - while functions may be outsourced, delegated or supported by AI systems, the regulated entity remains accountable for compliance with legal, regulatory and supervisory requirements;

    c) **proportionality**: the level of governance, risk management and oversight applied to AI should reflect the potential impact of the use case, taking into account the regulated entity's nature, scale and complexity;

    d) **integration with current frameworks** - existing governance, risk management, conduct, cyber, outsourcing, operational resilience, AML/ATF and recordkeeping frameworks remain primary;

    e) **robust lifecycle management** - risks should be considered from inception and selection through deployment, operation, monitoring, change and retirement;

    f) **responsible and ethical data use** - data used in AI systems should be subject to proportionate governance over provenance, permitted use, quality, confidentiality, security and potential inappropriate or unfair use;

    g) **explainability and transparency -** regulated entities should be able to understand, evidence and, where appropriate, explain the basis for material AI-supported decisions and the effectiveness of relevant controls; and

    h) **adaptability -** governance should evolve as AI capabilities, use cases, incidents and supervisory expectations develop.

44. The Authority does not propose a general requirement for firms to notify or obtain approval from the BMA before using AI. However, existing notification, approval, reporting or engagement requirements continue to apply where triggered by the underlying regulated activity, material outsourcing, business change, operational incident, model approval, customer disclosure or other applicable provision.

45. The Authority may engage with a regulated entity on its use of AI as part of its usual supervisory activities and in particular where an AI use case is material, affects critical operations or important business services, creates significant customer or market impact, involves substantial autonomy, relies on concentrated third parties, or otherwise raises prudential, conduct, cyber, AML/ATF or financial stability concerns.

46. Where a proposed AI use case is novel, complex or may raise uncertain regulatory questions, the Authority may consider engagement through its Innovation Hub or Sandbox, where appropriate. This reflects Bermuda's established approach of using the Innovation Hub and Sandbox to support

early supervisory dialogue for novel or complex use cases, while maintaining clear accountability within the applicable regulatory framework.

47. The Authority may use information gathered through supervisory engagement, Innovation Hub or Sandbox interactions, thematic reviews, market surveys and incident analysis to refine future guidance or consider targeted amendments where guidance proves insufficient.

IX. NEXT STEPS

48. The Authority invites comments on:
    a. whether the proposed GN appropriately explains how existing requirements apply to AI;
    b. whether the proposed materiality and risk assessment approach is workable across sectors; and
    c. whether further sector-specific examples would be useful.
49. Responses should be submitted by 30 October 2026 to policy@bma.bm
50. Following the consultation period, the Authority will assess responses received and refine the proposed GN as appropriate before issuing a final version. The Authority may also develop supplementary examples, supervisory tools or sector-specific materials where these would improve consistent implementation.
51. The Authority will issue updated guidance where material changes in the regulatory, technological or market landscape warrant doing so. It will continue to monitor international developments, AI-related incidents, market practices and emerging technologies, particularly GenAI, agentic AI and frontier models.

PART B - PROPOSED GUIDANCE NOTE

Guidance Note on the Responsible Use of Artificial Intelligence in Bermuda's Financial Services Sector

1. Purpose, status and interpretation

1. This GN explains how existing regulatory expectations may apply where a regulated entity develops, procures, deploys or uses AI in connection with a regulated activity or in a manner that may affect prudential soundness, customer or investor outcomes, market integrity, AML/ATF compliance, cybersecurity or operational resilience.
2. This GN is principles-based, technology-neutral and proportionate. It does not create new statutory or regulatory obligations, establish a separate AI licensing framework, prescribe a particular technology or model architecture, or require routine pre-approval of AI systems.
3. This GN should be read alongside the legislation, regulations, rules, codes, statements of principles, prudential standards, guidance and other regulatory requirements applicable to the regulated entity. Where an applicable legal or regulatory instrument imposes a specific requirement, that requirement prevails. This GN does not amend, replace or override any such instrument.
4. References in this GN to matters that a regulated entity 'should consider' describe the Authority's supervisory interpretation of how existing obligations and regulatory outcomes may be addressed in an AI context. They should not be read as creating a standalone obligation independent of the applicable legal and regulatory framework.
5. The considerations, examples and practices described in this GN are illustrative and non-exhaustive. A regulated entity may adopt alternative governance arrangements, risk management measures and controls where these are appropriate to the nature, scale, complexity, materiality and risk of the relevant AI use case and achieve outcomes consistent with its applicable obligations.
6. Responsibility for compliance remains with the regulated entity, including where it relies on an affiliate, group entity, third-party provider, open-source model, cloud platform or automated process.

2. Scope and application

7. The GN applies to AI used directly by a regulated entity and to AI embedded in products, services, systems or processes provided by third parties or group entities, where the use may affect regulated outcomes.
8. The GN covers traditional AI and machine learning, GenAI and agentic AI. Simpler rules-based or deterministic systems will not ordinarily be treated as AI, but equivalent governance may be appropriate where a complex quantitative method presents comparable risk.
9. The Authority considers the AI use case to be the appropriate primary level at which materiality and risk should ordinarily be assessed. The same model may support multiple use cases with different levels of materiality and risk. Conversely, a single use case may depend on several models, data sources, systems and providers that may need to be considered collectively.

10. For supervisory dialogue and internal governance purposes, regulated entities may find it useful to distinguish between:
    a. **Higher-impact AI use cases**, including those that affect regulated decision-making, financial outcomes, valuation, pricing, trading, underwriting or customer-facing decisions; and
    b. **Lower-impact AI use cases**, including internal productivity tools or operational support functions with limited direct impact on regulated outcomes.

This risk spectrum is illustrative and non-exhaustive. The appropriate level of governance, validation, oversight and documentation should be calibrated to the assessed materiality and risk of the use case.

11. The considerations and illustrative practices set out in this GN should be applied proportionately, having regard to the materiality and risk of the relevant AI use case. The relevant factors include the use case's impact, autonomy, complexity, explainability, data sensitivity, deployment context, scale, access to systems, customer or market effects, third-party dependencies, and relevance to critical operations or important business services.
12. Proportionality should operate principally at the level of the use case. For example, risks associated with a lower-risk internal productivity tool may be addressed through approved-use policies, access controls, user training and platform-level monitoring. A material decision-support tool may warrant documented ownership, testing, human review, monitoring and escalation. A higher-impact or more autonomous use case may warrant enhanced approval, more rigorous independent challenge, stronger auditability, tighter limits on autonomy, and more robust incident response arrangements and fallback measures. Where AI is provided by a group entity or third party, the assessment should also take account of the regulated entity's practical ability to obtain relevant information and the availability of compensating controls.
13. For the purposes of this GN, an AI use case may be considered 'material' where its failure or misuse could reasonably be expected to:
    a. affect customer, client, investor or policyholder outcomes;
    b. impact market integrity, pricing, valuation or trading decisions;
    c. influence access to financial services or risk selection;
    d. affect the continuity of important business services; or
    e. give rise to significant financial, prudential, legal or reputational risk.
14. The assessment of materiality should be based on the use case and its impact, rather than solely on the technical characteristics of the model or system. The assessment should ordinarily be undertaken through the entity's existing materiality, risk-classification or change-assessment methodologies under the applicable legal and supervisory framework, adapted where necessary to reflect relevant AI characteristics. This GN does not establish a separate AI-specific materiality test.

3. Existing regulatory frameworks remain primary

15. The Authority recognises that AI does not necessarily create entirely new categories of risk. AI may, however, introduce novel risk characteristics or amplify existing risks through increased autonomy, adaptiveness, opacity, interconnectedness, scale and speed. Relevant AI-specific risks may include hallucinations, model drift, automation bias, prompt injection, data poisoning, limited

explainability, concentration risk associated with common service providers and the use of agentic systems capable of autonomous decision-making or action.

16. AI-related risks may ordinarily be addressed through the regulated entity's existing governance, risk management and control arrangements. Depending on the use case, relevant arrangements may include those relating to corporate governance, prudential risk management, model risk, customer and market conduct, data governance, cybersecurity, outsourcing and third-party risk, operational resilience, AML/ATF and sanctions compliance, disclosure, recordkeeping, internal controls and audit.
17. A regulated entity is not expected to establish a separate AI governance framework solely because it uses AI, where its existing arrangements appropriately address the material risks of the relevant use cases. Where AI introduces materially different or heightened risks, the entity may need to enhance those arrangements. Depending on the nature, materiality and complexity of the use case, relevant enhancements may include clearer accountability, additional expertise, revised policies and controls, enhanced reporting or assurance, or dedicated governance arrangements.
18. For the purposes of this GN, the AI lifecycle may include, as applicable:
    a. identification and use-case definition;
    b. model or system selection;
    c. data preparation and development;
    d. validation and testing;
    e. deployment and integration;
    f. monitoring and performance management;
    g. change management; and
    h. retirement or decommissioning.
19. Relevant risks should be identified and managed at the applicable stages of the AI lifecycle, using the regulated entity's existing governance and risk management arrangements where appropriate.

Illustrative application process

The following diagram illustrates a process that regulated entities may use when applying this GN to an AI use case. It is intended to assist entities in identifying, assessing and managing AI-related risks through their existing governance and risk management arrangements. The process is illustrative and should be adapted to the nature, materiality and risk of the relevant use case.

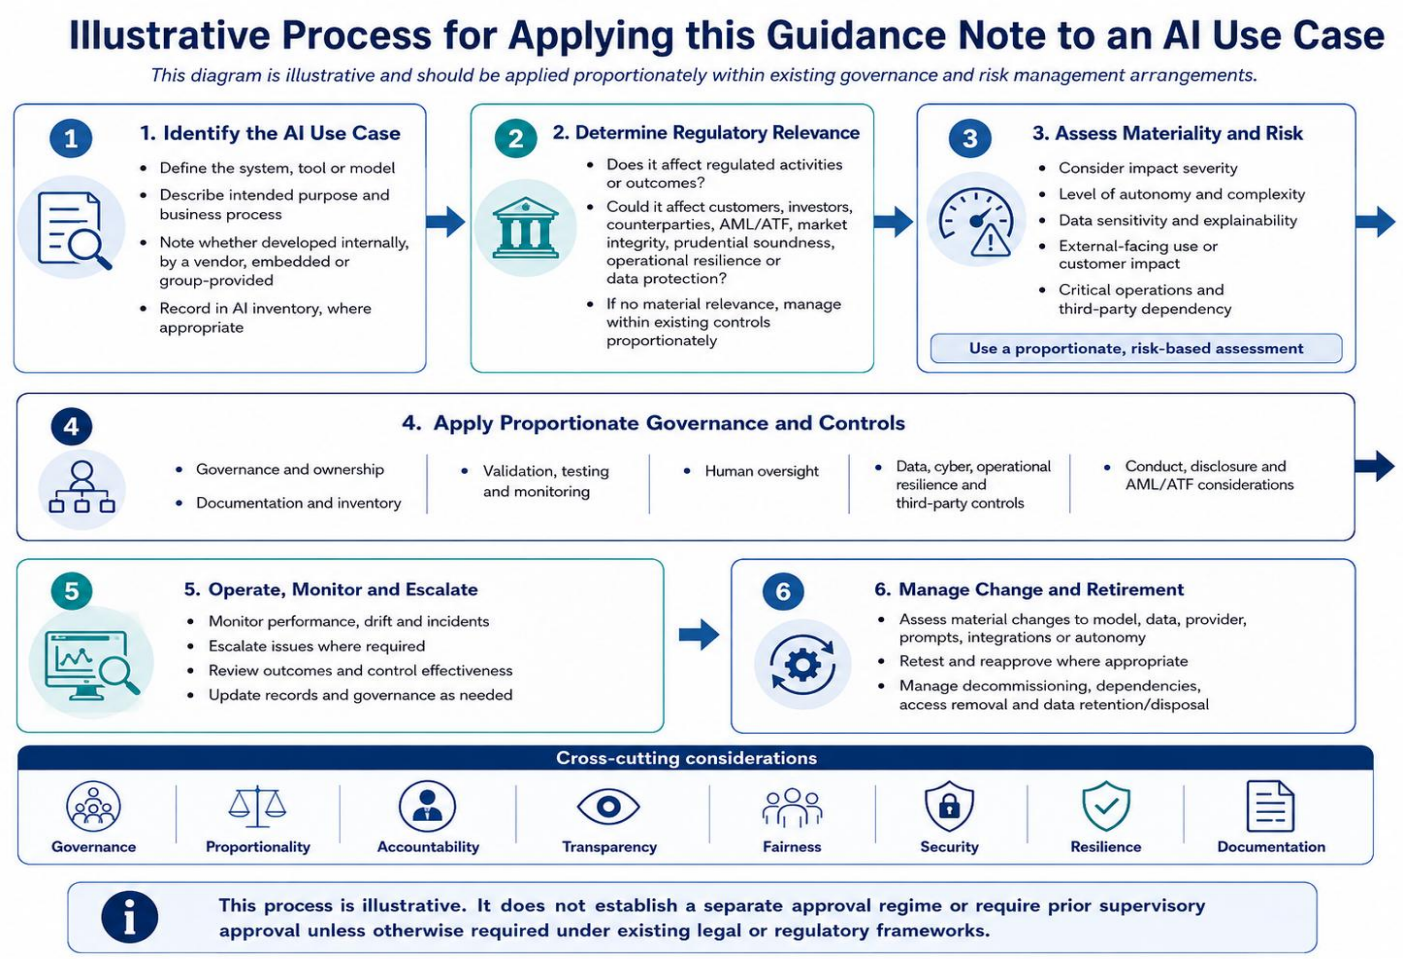

Figure 1: Illustrative process for applying this GN to an AI use case

4. Strategic direction, risk appetite and governance

20. The board and senior management should consider AI adoption in the context of the entity's business model, strategy, risk profile, risk appetite and applicable ethical standards and standards of conduct. Oversight should include both the potential benefits and the risks of adoption, including the implications of not adopting capabilities needed to manage fraud, cyber or other evolving risks.

21. The board should approve and oversee clear boundaries for AI adoption, proportionate to the entity and its use cases, including where appropriate clearly defined prohibited or restricted uses. These boundaries may include approval thresholds, restrictions on public GenAI tools, limitations on autonomous decisions or transactions, and requirements for human intervention in specified circumstances.

22. Prohibited uses may include, for example, use cases that:

    a. cannot be adequately controlled within risk appetite;

    b. would be inconsistent with applicable legal or regulatory obligations; or

    c. present unacceptable risks to customers, market integrity or the entity's safety and soundness.

23. The scope of such prohibitions should be proportionate to the entity's business model and risk profile.

24. Existing management information requirements should extend to information sufficient to oversee material AI. This may include material use cases, risk ratings, performance and drift indicators, incidents and near misses, third-party dependencies, audit findings, customer or market impacts, remediation and significant changes.

25. Board members and senior management should have an appropriate understanding of AI-related risks and use cases relevant to their responsibilities, enabling them to identify material opportunities and risks, challenge management, assess the effectiveness of controls and make informed decisions within their respective roles.

5. Governance roles, responsibilities and oversight

26. The rigour of the governance, oversight and controls applied to AI should be proportionate to the potential impact and risks arising from the AI application, considering the nature, scale and complexity of the regulated entity. Such measures may be embedded within existing governance and risk management frameworks.

27. The organisational arrangements used to govern AI should be proportionate to the nature, scale and complexity of the entity and the AI use case. Depending on the entity and the use case, appropriate outcomes may be achieved through existing board and management oversight, responsible officers, service providers, group arrangements, compliance, risk management, internal audit or other assurance mechanisms. Smaller or less complex entities may use simpler governance, documentation and review arrangements where these remain sufficient to identify the use case, assign accountability, understand material risks, challenge relevant outputs and demonstrate compliance with applicable obligations. A more extensive governance structure may be appropriate where the AI use case is material, highly autonomous and complex, dependent on sensitive data, customer- or investor-facing, or important to a regulated activity or critical operation.

28. Cross-functional coordination should involve relevant business, risk, compliance, legal, privacy, cyber, data, technology, model risk, operational resilience and internal audit expertise. The combination will depend on the AI use case and the nature, scale and complexity of the regulated financial entity.

29. Regulated entities should note that the use of AI does not alter the applicability of existing legal and regulatory obligations. Compliance should be assessed on an ongoing basis as AI use cases evolve.

6. AI identification, inventory and documentation

30. A regulated entity should maintain proportionate visibility over its AI use. This may be achieved through an AI inventory or through existing registers (such as model, system, application or third-party inventories), where these are sufficient to demonstrate effective oversight. Where an AI inventory, existing register or other suitable mechanism is used, the information recorded may include, as relevant and proportionate to the materiality and risk of the use case:
    a. identify the AI system and its purpose;
    b. assign accountable ownership within the entity;
    c. classify the use case based on its materiality and risk characteristics; and
    d. record key attributes relevant to oversight, including data sensitivity, level of autonomy, third-party dependencies and impact on regulated activities.

31. Any inventory, register or other mechanism used for this purpose should be maintained within the regulated entity's existing governance, assurance and recordkeeping arrangements, to the extent necessary to support effective oversight and compliance with applicable obligations.
32. For material or higher-risk AI use cases, information that may be relevant, where proportionate and necessary to support effective governance or evidence compliance with existing obligations, includes:
    a. the purpose and owner;
    b. intended, permitted and prohibited uses;
    c. model or system and rationale for selection;
    d. materiality and risk assessment;
    e. affected products, business lines and stakeholders;
    f. data sources, key assumptions and limitations;
    g. lifecycle status, dependencies and approvals;
    h. testing and monitoring results;
    i. incidents; and
    j. relevant third-party information.
33. Lower-risk productivity tools may be recorded at a category or platform level where individual use-case records would be disproportionate, provided the entity has adequate policies, approved-use boundaries, access controls and monitoring. Material use cases built on the same tool should be documented individually.
34. For GenAI and agentic AI systems, examples of information that may be relevant, where proportionate and necessary to support oversight, incident investigation or compliance with existing obligations, include:
    a. prompt or configuration versioning where relevant;
    b. tools and systems that an agent can access;
    c. permission levels, data sources used for retrieval;
    d. meaningful human approval points;
    e. activity logs, model or vendor version changes;
    f. known limitations; and,
    g. appropriate and effective override points and kill-switches for high-risk systems.
35. A regulated entity may consider maintaining information throughout the AI lifecycle where this supports its existing governance, internal assurance, incident investigation, recordkeeping or supervisory engagement obligations. This GN does not prescribe AI-specific records or a separate retention period. Existing recordkeeping and retention requirements apply only to the extent required by the applicable legal and regulatory framework and the nature of the underlying regulated activity, decision, process or incident.

36. Regulated entities should also distinguish between enterprise productivity tools used for general support functions and AI use cases that materially influence regulated activities. Controls for productivity tools may be applied at platform or category level, provided that boundaries on permitted use, data handling, and monitoring are clearly defined.

7. Materiality and risk assessment

37. A regulated entity should consider AI-related risks within its existing risk management framework. Where an AI use case is material or gives rise to heightened prudential, conduct, operational, cyber, AML/ATF, market integrity or customer risks, the entity should ensure that those risks are appropriately identified, assessed, managed and monitored throughout the AI lifecycle.

38. Materiality may include the importance of the AI use case to the entity's viability, critical operations, important business services and ability to meet its legal and regulatory obligations (see Section 2). Materiality may change as a use case scales, becomes embedded in a business process or begins to affect regulated activities, customers, investors, counterparties or other relevant stakeholders. A low-materiality use case may nonetheless present heightened risk – for example, an AI agent performing an administrative task may have broad access to sensitive data or systems. In such cases, materiality and risk should be reassessed following significant model, data, vendor, configuration or use changes, scaling to new business areas, incidents or near misses, material performance degradation or relevant external developments.

39. In assessing AI-related risk within its existing risk management framework, a regulated entity may have regard to:

    a. impact and severity (including effects on customers, investors, counterparties, market integrity, prudential soundness or important business services);

    b. autonomy and human oversight;

    c. complexity and explainability;

    d. data sensitivity and quality;

    e. deployment context and scale; and

    f. third-party, group or supply-chain dependencies.

40. Illustrative low- and enhanced-governance indicators for each dimension are set out in Annex C. A regulated entity may use its own internal classification or tiering approach; the Authority does not prescribe a particular taxonomy or methodology.

8. AI model or system selection

41. A regulated entity should consider whether the characteristics and limitations of an AI model or system are appropriate for the nature, materiality and risk of the intended use case and whether the entity can govern, validate, monitor and use the system consistently with its risk appetite and applicable legal and regulatory obligations.

42. Relevant considerations may include the system’s intended purpose, performance, explainability, robustness, data requirements, degree of autonomy, third-party dependencies, operational resilience and the potential impact on customers, counterparties, markets or critical operations.

43. For material or higher-risk use cases, the rationale for selecting the model, system, provider or configuration should be documented proportionately within the entity's existing governance, risk management, procurement or change-management records.

9. Data governance

44. A regulated entity should consider AI-related data risks within its existing data governance, privacy, confidentiality, security, outsourcing and recordkeeping arrangements. This includes, where relevant, data used for training, testing, retrieval, inference, monitoring and ongoing operation.
45. Relevant considerations may include data provenance, permitted use, lawful and ethical sourcing, sensitivity, quality, representativeness, timeliness, transformation, retention, privacy, security and cross-border implications. The depth of review should be proportionate to the materiality and risk of the use case and the extent to which the data may influence regulated activities, customer outcomes, prudential assessments or market integrity.
46. Ethical use of data should be considered through the regulated entity's existing governance, conduct, privacy, confidentiality, outsourcing, records management and risk management arrangements. This includes considering whether data used for AI has been obtained, accessed, shared and used in a manner consistent with applicable law, contractual rights, confidentiality obligations, customer or counterparty expectations, and the entity's own policies. The level of review should be proportionate to the materiality and risk of the use case, including whether the data may materially influence regulated decisions, customer or investor outcomes, prudential assessments, market activity or supervisory submissions.
47. Where a regulated entity relies on third-party AI systems or data that it cannot directly assess, it may use proportionate assurance, vendor due diligence, contractual protections, use restrictions, testing, output review or monitoring to manage residual risk.
48. For retrieval-augmented generation, memory-enabled or agentic systems, a regulated entity should consider whether retrieved or stored information could be inaccurate, stale, unauthorised, malicious or used outside its permitted context. Where material, existing data governance and security arrangements should be adapted to address those risks.

10. Validation, testing and performance management

49. Consistent with existing requirements for systems, controls and, where applicable, model risk management, a regulated entity should consider whether AI use cases require validation, testing or performance monitoring proportionate to their materiality and risk.
50. For material or higher-risk use cases, relevant considerations may include pre-deployment testing, review of intended use, performance monitoring, escalation arrangements and the ability to restrict or suspend use where outcomes appear unreliable, outside approved use or inconsistent with applicable obligations.
51. Testing and monitoring may consider, as relevant to the use case, accuracy, reliability, robustness, stability, bias, security, explainability, privacy, data leakage, performance degradation, model drift and the impact of changes in data, configuration, vendor models or operating conditions. Where a use case affects customers, investors or counterparties, testing and monitoring should also consider

whether outputs and decisions remain consistent with applicable market conduct considerations outlined in Section 19. Regulated entities may obtain relevant assurance by testing whether AI-enabled processes supporting critical customer services continue to deliver the outcomes required under the relevant contractual obligations.

52. For GenAI or other non-deterministic systems, testing may also consider hallucination risk, faithfulness to source material, prompt sensitivity, inappropriate content, data leakage and susceptibility to prompt injection. Where there is no single ground truth, a regulated entity may use a combination of sampled review, subject-matter expert assessment, benchmarking, red-teaming or other proportionate evaluation methods.

53. Records of validation, testing and monitoring should be maintained to the extent necessary to support the entity's existing governance, assurance, incident management, recordkeeping and supervisory engagement obligations. This GN does not prescribe a separate or universal AI model validation framework or a specific testing methodology.

11. Change management and retirement

54. A regulated entity should consider AI-related risks within its existing change management, technology risk, outsourcing, cyber and operational resilience arrangements. The level of review should reflect the materiality and risk of the use case.

55. For material or higher-risk use cases, relevant changes may include changes to the model, data, provider, configuration, prompt architecture, retrieval source, permissions, integration, intended use or level of autonomy. The entity should consider whether such changes require additional testing, approval, monitoring, stakeholder communication, fallback arrangements or other controls.

56. Where an AI system is retired or replaced, the entity should consider relevant dependencies, access removal, data retention or disposal, continuity of affected business processes, and the extent to which retired components remain embedded in other systems or workflows.

12. Meaningful human oversight

57. A regulated entity should consider the role of human oversight for AI systems within its existing governance and control framework, having regard to the materiality, risk, autonomy, complexity, explainability and potential consequences of the AI use case. References to human oversight in this GN illustrate one means by which a regulated entity may support compliance with existing governance, systems-and-controls, risk-management, conduct or prudential obligations. They do not create a standalone requirement for human review of every AI-supported output. The board and senior management, however, remain responsible for overseeing the adequacy of the entity's governance and control arrangements in accordance with the sectoral requirements applicable to the regulated entity.

58. Human oversight may take different forms, including human-in-the-loop review before execution, human-on-the-loop monitoring, exception-based intervention, periodic assurance, post-decision review, escalation triggers, transaction limits, suspension rights or other automated controls. The appropriate form should reflect the risk, autonomy and reversibility of the use case.

59. For use cases that produce or materially influence binding decisions, financial commitments, regulatory submissions, customer outcomes or important business services, the entity should consider whether human review is required before execution or whether post-decision monitoring, exception handling and escalation are sufficient.
60. Human oversight should be capable of providing meaningful challenge where it is relied upon as a control. This may include appropriate authority, competence, information, intervention time, escalation routes and awareness of automation bias or undue reliance on AI outputs.
61. Depending on the use case, oversight arrangements may include manual review, exception-based intervention, escalation, periodic monitoring, contestability or redress mechanisms, limits on autonomy, or the ability to restrict or suspend use. The form of oversight should be proportionate and should not be interpreted as requiring human review of every AI-supported output.
62. Where customers or other affected parties have a right under applicable law, contract or conduct requirements to challenge an outcome, the entity should consider how existing review, explanation, assistance or redress processes apply where AI materially influenced the outcome.

13. Explainability, transparency and fairness

63. Explainability, transparency and fairness should be considered in light of the entity's existing legal, regulatory, contractual, conduct, governance and risk management obligations. The appropriate level of explanation or disclosure will depend on those obligations and factors, such as the nature of the use case and outcome, the affected stakeholders and the purpose for which the information is required.
64. Explainability concerns the extent to which relevant persons can understand, challenge or evidence how an AI-supported output was generated or used. Transparency concerns the information made available to boards, senior management, control functions, auditors, supervisors, customers, investors, counterparties or other stakeholders.
65. Greater explainability or alternative controls may be appropriate where AI materially influences regulated activities, customer outcomes, regulatory submissions, financial commitments, market activity, or decisions that must be capable of audit, investigation, reconstruction or challenge.
66. Customer-facing disclosure is not required solely because AI is used within a process. Disclosure or explanation may be relevant where a person interacts directly with AI, AI materially influences an outcome, limitations of the AI system are material to informed use, or applicable law, contract or conduct requirements provide a right to explanation, review or redress.
67. Fairness should be assessed by reference to applicable legal, regulatory, conduct and contractual obligations. For consumer-facing use cases, this may include consideration of discriminatory, unfair or adverse outcomes. In wholesale or institutional markets, relevant considerations may include accuracy, consistency, conflicts, contractual transparency, market integrity and treatment aligned with negotiated terms.

14. Generative AI considerations

68. Where a regulated entity uses GenAI in connection with regulated activities, it should consider whether existing governance, data, cyber, conduct, outsourcing and operational resilience controls adequately address the characteristics of GenAI, including non-deterministic outputs and reliance on prompts, retrieved content or third-party models.
69. Relevant risks may include hallucinations, inaccurate or misleading content, inappropriate outputs, confidentiality breaches, data leakage, intellectual property issues, impersonation, prompt injection and over-reliance by users.
70. Depending on materiality and risk, proportionate controls may include approved-use boundaries, data-input restrictions, user guidance, output review, source verification, controlled retrieval sources, access controls, training, monitoring or restrictions on external distribution.
71. Content generated or materially supported by GenAI should be reviewed by appropriate personnel where errors could materially affect customers, investors, counterparties, supervisors, market integrity, legal compliance or the entity's regulated obligations.

15. Agentic AI considerations

72. Agentic AI may be capable of planning and executing multi-step tasks, invoking tools, accessing systems, communicating with other agents, retaining memory or taking actions with limited human intervention. These characteristics may create heightened risks where an agent can affect regulated activities, customer or counterparty outcomes, financial commitments, critical operations, important business services or sensitive systems.
73. A regulated entity considering agentic AI should assess whether the intended outcome could be achieved through an alternative process with materially lower risk. Where agentic capability is used, the entity may consider proportionate limits on autonomy, objectives, duration, tools, data access, permissions and ability to act, and where appropriate, design it so that agentic AI systems operate within constrained or isolated environments, with bounded permissions to limit unintended actions and the propagation of errors, and ensure that they are adequately tested before being scaled.
74. Where an agent can invoke tools, APIs, databases, code-execution environments, communication channels, payment functions or other systems, the entity should consider controls over permitted tools, permitted purposes, approval thresholds, access rights and monitoring. For higher-impact actions, controls should not rely solely on prompts or model instructions.
75. For material or higher-risk agentic use cases, the entity should consider whether records of relevant actions, tool calls, access decisions, approvals, exceptions and outcomes are needed to support oversight, investigation, incident management, auditability or contestability under existing obligations.
76. Existing cyber, operational resilience, outsourcing, access management and incident management frameworks should be applied to agentic AI where relevant. This may include monitoring for unusual access or tool-use patterns, scope creep, repeated failures, compromised instructions, cross-agent trust issues or propagation of errors through connected workflows.

16. Cybersecurity and ICT risk

77. A regulated entity should consider AI-related risks within its existing cyber risk management framework including applicable asset identification, threat intelligence, vulnerability management, access management, incident response and third-party risk arrangements. Relevant risks may

include evasion, data or model poisoning, prompt injection, jailbreaking, model or data extraction, sensitive information disclosure, deepfakes, AI-enabled phishing, malicious code generation, memory or context poisoning, tool misuse, identity and privilege abuse, insecure agent communications, supply-chain compromise and unauthorised use of AI infrastructure.

78. The nature and extent of controls should be proportionate to the materiality and risk of the AI use case. Depending on the use case, existing cyber controls may need to be adapted to address AI-specific risks, including secure configuration, access controls, least privilege, monitoring, logging, change management, data loss prevention, containment, recovery and controls over connected tools or agents.
79. A regulated entity should consider how unauthorised or unmanaged AI use, including shadow AI, embedded AI features and unapproved uses of approved platforms, is identified and managed through existing technology, information security, data governance, training and acceptable-use arrangements.
80. Where relevant to material or higher-risk use cases, AI-related attack scenarios may be considered within existing vulnerability testing, red-teaming, tabletop exercises or operational resilience exercises. Such scenarios may include attacks against AI systems, AI-enabled attacks against the entity or its customers, and propagation through connected tools, agents, data sources or service providers.
81. Established technical references, including but not limited to the National Institute of Standards and Technology (NIST), the MITRE Adversarial Threat Landscape for Artificial-Intelligence Systems (MITRE ATLAS) and the Open Worldwide Application Security Project (OWASP) materials, may assist entities in developing threat models or testing approaches. These references are illustrative only and do not replace applicable legal or regulatory requirements, nor do they provide a safe harbour or guarantee compliance.

17. Operational resilience and incident management

82. Where AI materially supports an important business service, critical operation or material function, a regulated entity should consider relevant AI dependencies and failure modes within its existing operational resilience, business continuity, incident management and outsourcing arrangements.
83. Relevant considerations may include dependence on third-party AI services, model degradation, corrupted retrieval or memory data, inaccurate or misleading outputs, data leakage, deepfake-enabled fraud, unauthorised agent actions, compromised tools or identities, cascading failures across connected systems or common-provider disruption.
84. Depending on the materiality and risk of the use case, continuity or resilience arrangements may include manual processing, alternative models or providers, restricted functionality, isolation of affected components, activity or transaction limits, rollback, data portability or other recovery options. The feasibility and cost of such arrangements may be considered when selecting, contracting for or materially changing AI systems.
85. AI-related incidents should be managed through existing incident management frameworks. Root-cause analysis may consider AI-specific factors such as data drift, model or vendor updates,

emergent behaviour, human over-reliance, changes in prompts or retrieval sources, and interaction between connected systems.

86. Existing regulatory reporting or notification requirements apply where an AI-related incident meets the relevant threshold. This GN does not create a separate universal AI incident reporting requirement.

18. Third-party, group and supply-chain risk

87. A regulated entity should consider AI-related risks within its existing outsourcing, third-party risk management, group governance, operational resilience and cyber risk frameworks. The entity remains responsible for understanding and managing the risks of AI use in connection with its regulated activities, including where AI is provided by a group entity, vendor, embedded software feature, open-source component or external platform.

88. The depth of due diligence and ongoing oversight should be proportionate to the materiality and risk of the AI use case, the nature of the provider relationship, the entity's reliance on the system and the extent to which the AI use case may affect regulated activities, customers, investors, counterparties, critical operations, important business services or market integrity.

89. Relevant due diligence considerations may include the provider's governance, experience, resilience, data practices, security, model performance, limitations, explainability, incident history, subcontractors, dependencies and change management arrangements. The relevant AI supply chain may include models, datasets, software libraries, tools, plugins, APIs, agent frameworks, update channels, cloud infrastructure, hardware infrastructure or other components that could materially affect performance, security or continuity.

90. Contracts and assurance arrangements may, where proportionate and reasonably achievable, address permitted use, service levels, data access and security, material model or feature changes, incident notification, audit or assurance rights, regulatory access, subcontracting, continuity, portability, termination and exit. Standardised disclosures, model or system cards, certifications, independent assurance, pooled audits or proportionate dependency inventories may support due diligence but are not mandatory forms of evidence.

91. The Authority recognises that regulated entities may not always have access to source code, model weights, training data or full technical documentation, particularly where AI is provided through global platforms, embedded software or group systems.

92. Where information or contractual rights are limited, a regulated entity may consider compensating controls such as use restrictions, additional testing, output review, benchmarking, enhanced monitoring, staged deployment, warranties, indemnities or limiting the system to lower-risk use cases. The objective is to obtain sufficient assurance to manage risk within appetite, rather than achieve full technical transparency.

93. Embedded AI features in existing products should be considered within proportionate governance arrangements, particularly where new or materially changed AI capabilities may alter the risk profile of an existing service, system or workflow.

94. A regulated entity should consider material concentration risks arising from reliance on common providers, group platforms, cloud infrastructure, hardware, foundation models, data providers, agent

platforms or vertically integrated ecosystems. Where relevant, this may inform continuity planning, scenario testing, alternative provider assessment, data portability, exit planning or supervisory engagement.

95. Where an AI system is developed, governed or operated at group level, the BMA-regulated entity should retain sufficient local oversight to understand the purpose, limitations and risk profile of the AI use case and to assess whether its use is consistent with the entity's regulatory obligations, risk appetite and business context.

96. Reliance on group frameworks does not displace the responsibility of the Bermuda board and senior management to oversee AI use within the regulated entity. However, the entity may consider adopting and applying controls and measures drawn from the group where appropriate for the Bermuda entity's business model, risk profile and regulatory obligations.

97. Regulated entities are not expected to obtain proprietary model information where this is not reasonably available or contractually supported. Where full technical transparency is unavailable, entities should focus on whether the available information, assurance and controls are sufficient to support responsible use within existing governance and risk management arrangements.

19. Conduct of business, market integrity and disclosure

98. Consistent with the validation and monitoring expectations in Section 10, a regulated entity should assess and manage AI use within its existing conduct of business obligations, including disclosure, suitability, fair treatment, conflicts of interest, market conduct, marketing, complaints handling, and customer or investor protection obligations. The method used to make, support or communicate a customer-, investor- or counterparty-related decision does not change the underlying obligation owed to customers, investors, counterparties or other affected stakeholders.

99. Communications, advice, disclosures, customer interactions or counterparty communications generated or materially supported by an AI system should meet the same applicable standards as communications delivered by personnel or other non-AI systems. Where an AI system interacts with customers, investors, counterparties or other external stakeholders, the regulated entity remains responsible for ensuring that communications are accurate, fair, clear, not misleading, and consistent with applicable legal, regulatory and contractual obligations. The entity should also consider whether relevant notices, explanations, user prompts, disclosures of material limitations or other information are appropriate to the nature of the interaction and necessary to meet applicable conduct, disclosure, fair treatment, contractual or customer and investor protection requirements. Where appropriate to the nature and significance of the interaction, the entity should also consider whether customers or other affected persons should have a reasonable means of escalating the matter to, or obtaining assistance from, an appropriately authorised human representative.

100. Claims about the entity's use, capability, performance or benefits of AI should be accurate, not misleading, and supported by documentation. Exaggerating, concealing or otherwise misrepresenting the use of AI, including through 'AI washing', may create conduct, disclosure, governance and reputational risks, particularly where customers, investors or counterparties could be led to make decisions on an inaccurate understanding of the entity's technology, controls or capabilities.

101. Where AI is used for trading, investment, pricing, market surveillance, portfolio management or other market-facing activities, a regulated entity should consider whether AI use could create or amplify market integrity risks, including correlated behaviour, herding, procyclicality, volatility, conflicts of interest, inappropriate reliance on common models or data, or ineffective challenge of automated outputs. Where such risks are material, the entity should address them through proportionate controls, monitoring, limits, escalation arrangements, human review, independent challenge or other mitigants consistent with its existing market integrity and risk management obligations.

20. AML/ATF and Sanctions

102. The Authority recognises that AI may enhance AML/ATF and KYC processes, sanctions screening, fraud detection, adverse media review, beneficial ownership analysis and investigations. Where applicable, the regulated entity remains responsible for the effectiveness of its systems and compliance with AML/ATF and sanctions requirements.

103. In using AI for financial crime controls, the regulated entity should consider validation and monitoring proportionate to the use case, including false positives, false negatives, drift, data quality, explainability sufficient for investigation and reporting, automation bias in alert handling, potential inappropriate or discriminatory outcomes and the effect of vendor, model or data changes. Where relevant, the entity should consider whether particular customer groups, jurisdictions, business models or other categories are subject to materially different outcomes that cannot be justified by the underlying financial crime risk.

104. AI systems do not replace Money Laundering Reporting Officer accountability or other accountable decision-making under applicable AML/ATF and sanctions requirements. Where GenAI is used to generate investigative narratives, sanctions analysis or draft regulatory reports, outputs should be verified for accuracy, completeness, evidential support and appropriate handling of confidential information before reliance or submission.

105. Where relevant, controls should address emerging risks such as synthetic identity, deepfakes, AI-enabled fraud, evasion of transaction monitoring, sanctions screening limitations and the auditability of decisions to close or escalate alerts.

21. Organisational adaptability and capability

106. A regulated entity should consider whether it has internal capability proportionate to the nature, scale, materiality and risk of its AI adoption. This may include appropriate AI literacy and role-specific training for relevant employees, officers, senior management, control functions and board members, in alignment with existing requirements for employee and board-level training.

107. Training may address approved uses, limitations, automation bias, data handling, cyber threats, escalation routes, incident lessons and the entity's internal policies for AI use. The scope and frequency of training should be proportionate to the roles involved and the materiality and risk of the relevant AI use cases.

108. The entity should consider relevant external developments, including supervisory publications, incidents, technology changes, threats, industry practices and legal requirements. Lessons learned from incidents, near misses, testing, drift monitoring, assurance reviews or

supervisory engagement should be incorporated into policies, controls, communications and training, as appropriate.

109. Smaller entities may use external or group resources and arrangements where appropriate, provided that the regulated entity retains sufficient understanding, accountability, effective challenge and access to relevant information to oversee AI use within its business.

22. Recordkeeping and auditability

110. A regulated entity should consider whether its existing recordkeeping, logging, audit trail and incident management arrangements are sufficient to support responsible AI use, having regard to the materiality and risk of the relevant use case.

111. For material or higher-risk AI use cases, and as appropriate under existing obligations, relevant records may include information sufficient to understand system behaviour, investigate incidents or unexpected outcomes, support internal assurance, evidence escalation or reconstruct material decisions including, where relevant, information on material inputs and outputs, human review, overrides, reliance and the basis for the final decision or action.

112. Records and logs, where required or maintained for these purposes, should be managed within the entity's existing recordkeeping and retention arrangements and in accordance with applicable requirements.

113. AI-related incidents should be managed as set out in Section 17, having particular regard to impacts on customers, investors, counterparties, markets and operations, as well as applicable reporting or escalation requirements.

23. Supervisory engagement and documentation

114. The Authority will assess AI use through existing supervisory processes. The nature and depth of engagement will be proportionate to the use case and the applicable regulatory framework.

115. In this regard, information the Authority may request can include governance documents, risk appetite and tolerances, board or committee documents and reporting, inventories, materiality and risk assessments, selection rationale, data governance records, validation and testing results, monitoring and drift reports, human oversight arrangements, incident records, third-party due diligence, contracts, assurance reports, continuity plans and internal audit findings.

116. Documentation should be proportionate and may be maintained within existing systems and records. Where information and documentation are held by a group entity or third party, the BMA-regulated entity should have sufficient access to support its responsibilities and supervisory engagement as may be required by the Authority.

117. The Authority may request additional information and documentation where an AI use case presents material prudential, conduct, market integrity, cyber, AML/ATF, operational resilience or financial stability concerns.

118. For material or higher-risk AI use cases, the Authority may assess whether the information available through the regulated entity's existing governance, risk management, control and recordkeeping arrangements is sufficient to demonstrate compliance with the applicable legal and regulatory framework.

ANNEX A – GLOSSARY OF TERMS

The following terms are used in this Guidance Note. The descriptions are intended to support consistent understanding and do not create standalone legal definitions. They should be read in context and subject to the applicable statutory and regulatory framework.

In particular, the Guidance Note distinguishes between traditional AI and machine learning (which generally identify patterns or make predictions from data), Generative AI (which creates new content, often in response to prompts) and Agentic AI (which can plan, reason and execute multi-step tasks with varying degrees of autonomy). See the respective entries below for further details. These distinctions matter because the nature of the system may affect the appropriate governance, testing, human oversight, cyber controls and operational safeguards.

Term	Description
Adversarial attack	An attempt to manipulate an AI system, its data, model, inputs, connected tools or operating environment in order to cause unintended behaviour, disclosure, disruption or other harm.
Agentic AI	An AI system or configuration capable of planning or sequencing multi-step tasks, invoking tools or systems, retaining or using context, and taking actions with varying degrees of autonomy.
AI	A machine-based or software-based system that, for explicit or implicit objectives, processes inputs to generate outputs such as predictions, content, recommendations, decisions or actions that may influence digital, operational or business environments, with varying degrees of autonomy or human oversight. The term does not extend to conventional calculators, static spreadsheets, simple rules-based tools or longstanding quantitative methods merely because they process data or produce an output; the substance, context, autonomy, adaptiveness and risk of the use case determine whether AI governance considerations are relevant. Conversely, a system not formally labelled as 'AI' may still warrant AI governance considerations where it demonstrates comparable autonomous or adaptive characteristics. This GN uses 'AI' and 'AI system' interchangeably.
AI inventory	A proportionate record or existing register through which an entity maintains visibility over material AI systems or use cases, including relevant ownership, purpose, risk, lifecycle, data and dependency information.
AI model	A computational component of an AI system that processes inputs to generate outputs, including predictions, classifications, recommendations, content, scores or other results.
AI-related incident	An event involving the development, deployment or use of an AI system that causes, or has the potential to cause, material harm, operational disruption, regulatory non-compliance, financial loss, inaccurate or unfair outcomes, data compromise or other adverse consequences.
AI system	See 'AI'.
AI use case	The specific business purpose and operational context in which an AI application, model, system, tool or capability is used.
AI washing	Exaggerating, misstating or making unsupported claims about an entity's use, capability, performance or benefits of AI.
Assurance review	An independent or appropriately objective assessment of whether an AI system and its related governance, risk management and controls are appropriately designed and operating effectively. Depending on the nature and materiality of the use case, this may be performed internally or by a suitably qualified external party.
Automation bias	The tendency to over-rely on automated outputs or recommendations without sufficient independent challenge.
Cascading failure	A failure or error that propagates through connected models, agents, tools, systems or processes and causes wider disruption or harm.

Cross-agent trust issues	Risks arising where one AI agent relies on another agent's identity, authority, instructions, data or outputs without adequate verification, potentially enabling errors, manipulation, unauthorised actions or compromised instructions to propagate across connected systems or workflows.
Data drift	A change over time in the characteristics or distribution of data used by an AI system, which may affect performance.
Data poisoning	The manipulation or corruption of data used to train, fine-tune, retrieve information for, or otherwise operate an AI system.
Deepfake	Synthetic or manipulated image, video, audio or other media generated or altered using AI or related tools in a manner that may falsely appear to depict a real person, event, communication or instruction.
Drift monitoring	The ongoing or periodic monitoring of an AI system to identify material changes in its data, performance, behaviour or outputs over time that may affect its accuracy, reliability, fairness or continued suitability for its intended purpose.
Embedded AI	AI functionality integrated into a broader product, platform, service, software application or business process, whether or not separately labelled as AI.
Emergent behaviour	Behaviour, capabilities or outcomes exhibited by an AI system that were not explicitly designed, programmed or reasonably anticipated, and which may arise from its operation, adaptation or interaction with users, tools, data or other systems.
Evasion	An attempt to craft inputs or behaviour to avoid detection, classification, monitoring or control by an AI system.
Explainability	The extent to which relevant persons can understand, challenge or evidence how an AI-supported output was generated or used.
Foundation model	A model trained on broad data at scale that can be adapted or used for a wide range of downstream tasks.
Generative AI (GenAI)	AI that generates new content, including text, images, audio, video or code, often in response to user prompts.
Group-provided AI	An AI model, system, application, component or service developed, hosted, maintained, governed or materially supported by another entity within the regulated entity's group.
Hallucination	An output generated by an AI system that appears plausible but is inaccurate, unsupported or fabricated.
Human over-reliance	Excessive or inappropriate reliance on an AI system's outputs, recommendations or actions without sufficient human scrutiny, independent judgment or verification, including where users defer to the system despite indications that its output may be inaccurate or unsuitable.
Human oversight	Meaningful human involvement that may enable an appropriately placed person with relevant information, competence and authority to review, challenge, intervene in or remedy AI-enabled outcomes.

Jailbreaking	An attempt to bypass or weaken restrictions, safeguards or intended-use boundaries in an AI system.
Machine learning	A method by which a system learns patterns or relationships from data to make predictions, recommendations or decisions without each rule being explicitly programmed.
Materiality	The importance of an AI use case to an entity's viability, critical operations, important business services, core business lines, regulated activities or ability to meet legal and regulatory obligations.
Model drift	A deterioration or change in model performance or behaviour over time as underlying relationships or conditions change.
Model or system card	A standardised summary of relevant information about an AI model or system, which may include intended use, limitations, performance, data, governance or risk information.
Model poisoning	The malicious alteration of an AI model, its parameters or related components to affect its behaviour or outputs.
Near miss	An event involving an AI system that could reasonably have resulted in an incident but did not, whether because it was detected and addressed in time, a control operated effectively, or the adverse outcome was otherwise avoided.
Non-deterministic output	An output that may vary across repeated uses of the same or similar inputs, including because of model design, configuration, sampling, context or system updates.
Prompt	An instruction, query, context or other input provided to a generative or agentic AI system to influence its output or action.
Prompt injection	An attempt to manipulate a generative or agentic AI system through instructions embedded directly or indirectly in prompts, documents, webpages, messages or other content.
Retrieval-augmented generation (RAG)	A technique that supplements a generative AI model with information retrieved from selected external or internal data sources.
Retrieval source	An internal or external source of information used by a retrieval-augmented, memory-enabled or agentic AI system to generate or support outputs.
Scope creep	The gradual or unauthorised expansion of an agentic AI system's activities, permissions, objectives or use of tools beyond its approved mandate, intended purpose or established operating limits.
Shadow AI	AI tools, features or uses adopted without appropriate authorisation, visibility or governance by the entity.
System instruction	A higher-priority instruction or configuration setting that guides the behaviour, boundaries or role of a generative or agentic AI system.
Third-party AI	An AI model, system, application, component or service developed, provided, hosted, maintained or materially supported by a third-party service provider.

Transparency	The extent to which appropriate information about an AI system, its purpose, use, limitations, governance or role in a process is made available to relevant stakeholders, having regard to the use case and applicable obligations.

ANNEX B - AI RISK THEMES AND EXISTING BERMUDA REGULATORY FRAMEWORKS

The table below illustrates how common AI risk themes may relate to existing Bermuda regulatory framework areas. It is included for orientation only and does not replace the applicable legislation, rules, codes, prudential standards, statements of principles, guidance or supervisory requirements relevant to each regulated entity.

AI risk theme	Existing Bermuda framework area	What this GN clarifies
Governance and accountability	Corporate governance, board oversight, senior management responsibility, statements of principles, codes of conduct and sectoral prudential frameworks	AI can be governed through existing accountability, risk appetite, delegation, management information and effective challenge arrangements.
Risk management and internal controls	Enterprise risk management, systems and controls, internal audit, compliance and assurance requirements	AI risks should be identified, assessed, managed, monitored and assured proportionately to the materiality and risk of the use case.
Model, system and performance risk	Prudential risk management, actuarial, valuation, pricing, trading, underwriting, reserving and model governance practices where applicable	Material AI use cases should be subject to proportionate selection, validation, testing, monitoring, drift management and change control.
Data governance and confidentiality	Books and records, confidentiality, privacy, data security, outsourcing and operational controls	AI data should be assessed for provenance, permitted use, lawful and ethical sourcing, quality, sensitivity, security, retention and cross-border considerations.
Cybersecurity and ICT risk	Cyber risk management, technology risk, access management, vulnerability management and incident response frameworks	AI-specific threats such as prompt injection, data poisoning, model extraction, deepfakes, insecure agents and shadow AI should be considered within existing cyber arrangements.
Operational resilience	Operational resilience, business continuity, important business services, incident management and dependency management	AI dependencies and failure modes should be considered where AI supports critical operations, material functions or important business services.
Outsourcing, third-party and group risk	Outsourcing, third-party risk management, group governance, service provider oversight and business continuity frameworks	Reliance on vendors, group platforms, embedded AI, cloud services, foundation models and AI supply chains does not transfer accountability from the Bermuda regulated entity.
Conduct, disclosure and market integrity	Conduct of business, fair treatment, disclosure, conflicts, suitability, complaints, market conduct and investor/customer protection requirements	AI-supported communications, decisions and market-facing activities should meet the same standards as non-AI processes, including controls against misleading claims or AI washing.

AI risk theme	Existing Bermuda framework area	What this GN clarifies
AML/ATF and sanctions	AML/ATF, sanctions, KYC, transaction monitoring, suspicious activity reporting and MLRO accountability frameworks	AI may support financial crime controls, but regulated entities remain responsible for effectiveness, explainability, validation, auditability and accountable decision-making.
Supervisory evidence and engagement	Existing supervisory powers, onsite/offsite supervision, reporting, notifications, thematic reviews and information requests	The Authority may request proportionate evidence of AI governance, risk assessment, validation, monitoring, incidents, third-party assurance and board oversight where relevant.

ANNEX C - ILLUSTRATIVE PROPORTIONALITY INDICATORS

The indicators below are illustrative only. They are not cumulative, exhaustive or determinative. A regulated entity may use different indicators or classifications within its existing risk management framework, provided the approach is proportionate to the materiality and risk of the AI use case.

Dimension	Lower-intensity indicators	Enhanced-governance indicators
Impact	Internal support; readily reversible; limited financial, customer, investor, counterparty or market effect.	Binding or high-impact decisions; prudential, customer, investor, counterparty or market consequences.
Autonomy	Recommendation only; frequent human review; no independent system access or action.	Autonomous actions; broad tool, data or system access; limited real-time oversight.
Complexity and explainability	Transparent logic; stable model; well-understood assumptions and limitations.	Opaque model; foundation model dependency; dynamic updates; difficult validation, challenge or explanation.
Data	Public or low-sensitivity data; narrow dataset; limited transformation or enrichment.	Personal, confidential, regulated, high-volume, cross-border, third-party or sensitive data.
Scale and criticality	Small pilot; non-critical process; limited users; limited transaction or decision volume.	Critical operations, important business services, broad deployment, high transaction volume or significant operational dependency.
External-facing use and stakeholder impact	Internal use only; no direct customer, investor, counterparty or market communication.	Customer, investor, counterparty or market-facing use; AI materially influences communications, outcomes, disclosures or regulated decisions.
Third-party, group or supply-chain dependency	Substitutable provider or group service; adequate assurance; manageable exit or fallback options.	Concentrated provider, group platform or embedded AI; limited assurance, portability, visibility or exit options.

Dimension	Lower-intensity indicators	Enhanced-governance indicators
Changeability and monitoring	Stable configuration; infrequent updates; changes controlled through normal technology processes.	Dynamic model or vendor updates; changing retrieval sources; frequent prompt or configuration changes; limited visibility over updates.
Human oversight	Clear human review, challenge and escalation; outputs used as support only.	Limited ability to intervene; high reliance on automated outputs; oversight may be impractical, delayed or ineffective.
Legal, regulatory or conduct sensitivity	Limited effect on regulated activities or legal obligations; no material customer or market impact.	Use case materially affects regulatory submissions, customer outcomes, investor interests, AML/ATF, market integrity, disclosure, suitability, claims, underwriting, pricing or other regulated obligations.

ANNEX D - ILLUSTRATIVE SECTOR USE CASES

The examples below are illustrative only. They are not exhaustive, determinative or intended to create sector-specific rules. The appropriate governance, documentation, validation, monitoring and oversight should be proportionate to the materiality and risk of the AI use case, having regard to the entity's applicable legal and regulatory obligations, business model, affected stakeholders, data sensitivity, level of autonomy, third-party dependencies and deployment context.

A use case listed in this annex should not automatically be treated as high-risk or low-risk. The same AI tool may present different risks depending on how it is configured, the data it uses, the level of autonomy permitted, the affected stakeholders, the extent of human oversight and the consequences of error, misuse or failure.

For illustrative purposes, an AI tool deployed to draft internal meeting summaries generally presents a lower regulatory risk profile, whereas a highly autonomous model executing algorithmic trades or denying insurance claims without meaningful human intervention would warrant significantly higher supervisory scrutiny.

Sector / function	Illustrative AI use cases	Possible heightened considerations
Insurance and reinsurance	Underwriting support; pricing; reserving; claims triage; fraud detection; catastrophe and exposure modelling; capital modelling; portfolio optimisation; document extraction.	Accuracy and explainability of outputs; data quality and representativeness; actuarial and underwriting judgement; treatment of policyholders, cedants, beneficiaries and counterparties; claims handling standards; model validation; auditability of material assumptions; reliance on third-party data, models or platforms.
Funds and investment business	AI-supported trading; portfolio analytics; investment research; NAV or fair valuation support; liquidity and redemption modelling; investor reporting; compliance surveillance; market abuse monitoring.	Market integrity; conflicts of interest; explainability of investment or valuation outputs; human oversight of trading or allocation decisions; data provenance; valuation governance; liquidity risk; investor disclosure; recordkeeping; third-party model or data-provider reliance.

Sector / function	Illustrative AI use cases	Possible heightened considerations
Digital assets and payments	Blockchain analytics; transaction monitoring; sanctions screening; fraud detection; smart-contract review; payment routing; cyber monitoring; wallet or transaction risk scoring.	AML/ATF and sanctions effectiveness; false positives and false negatives; auditability of alert handling; cyber resilience; model drift; data-source limitations; explainability of risk scores; smart-contract vulnerabilities; reliance on vendors, APIs, infrastructure providers or analytics platforms.
Banking and money service business	Credit decision support; AML/KYC; fraud monitoring; customer service tools; transaction monitoring; treasury and liquidity analysis; operational workflow automation.	Fair treatment and customer impact; explainability of customer-facing or credit-related outcomes; human review and escalation; data privacy and confidentiality; AML/ATF effectiveness; fraud typology changes; resilience of high-volume processes; outsourcing and vendor assurance.
Trust, corporate service provider and fund administration	Beneficial ownership mapping; due diligence; adverse media screening; document review; regulatory and tax reporting support; client onboarding; investor servicing; workflow automation.	Accuracy and completeness of due diligence outputs; explainability of screening results; confidentiality and data handling; beneficial ownership evidence; audit trail; human review of material findings; regulatory filing accuracy; reliance on third-party data sources or document-analysis tools.
Cross-sector risk, operations and governance	GenAI productivity tools; coding assistance; knowledge retrieval; policy and procedure drafting; cyber defence; incident analysis; regulatory change monitoring; board or management reporting support.	Approved-use boundaries; data-input restrictions; confidentiality; hallucination and source-verification risk; prompt injection; access controls; output review before external use; records of material reliance; monitoring of shadow AI; training and escalation arrangements.

ANNEX E - INTERNATIONAL AND EXTERNAL SOURCES

The sources below informed the Authority's benchmarking and policy analysis. They are included for context only. Inclusion does not mean that the Authority has adopted any source as binding guidance or incorporated every expectation, standard or control into this GN.

1) Financial Stability Board, Sound Practices for Responsible Adoption of Artificial Intelligence
2) Financial Stability Board, The Financial Stability Implications of Artificial Intelligence
3) International Association of Insurance Supervisors, Application Paper on the Supervision of Artificial Intelligence
4) International Organization of Securities Commissions, Supervisory Toolkit for AI Use in Capital Markets: Standalone Toolkit
5) National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0)
6) National Association of Insurance Commissioners, Model Bulletin: Use of Artificial Intelligence Systems by Insurers
7) Prudential Regulation Authority, SS1/23 Model Risk Management Principles for Banks
8) Financial Conduct Authority, AI Update and related publications
9) European Union, Regulation (EU) 2024/1689 (Artificial Intelligence Act) and related definitional guidance
10) US Department of the Treasury, Artificial Intelligence in Financial Services: Uses, Opportunities and Risks
11) Federal Financial Supervisory Authority (BaFin), Guidance on ICT Risks in the Use of AI at Financial Entities
12) Monetary Authority of Singapore and MindForge materials on operationalising AI risk management
13) National Institute of Standards and Technology, Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations (NIST AI 100-2e2025)
14) MITRE, Adversarial Threat Landscape for Artificial-Intelligence Systems (ATLAS)
15) OWASP GenAI Security Project, OWASP Top 10 for LLM Applications 2025
16) OWASP GenAI Security Project, OWASP Top 10 for Agentic Applications 2026 and related public Agentic Security Initiative guidance

ANNEX F - VISUAL SUMMARY OF THE PROPOSED FRAMEWORK