



Laurenceo

INDEPENDENT RESEARCH ON AI GOVERNANCE

PUBLIC CONSULTATION RESPONSE

Evidence by Default: Operationalising the Responsible Use of AI in Bermuda's Financial Services Sector

A response to the Bermuda Monetary Authority Consultation Paper,
"Guidance Note: The Responsible Use of Artificial Intelligence in
Bermuda's Financial Services Sector"

SUBMITTED TO

Bermuda Monetary Authority, Policy
policy@bma.bm

CONSULTATION ISSUED

14 August 2026

RESPONDENT

Laurenceo, Hamilton, Bermuda
laurenceo.com

FEEDBACK DEADLINE

30 October 2026

PREPARED BY

Jorai Laurenceo, Founder

REFERENCE

LAU-BMA-AI-2026-01 | October 2026

Executive Summary and Institutional Positioning

Laurenceo welcomes the Authority's Consultation Paper and supports its central design choice: to explain how existing legal, regulatory and supervisory obligations apply to artificial intelligence, rather than to create a stand-alone AI licensing or approval regime. A principles-based, technology-neutral and proportionate approach is the right foundation for a jurisdiction whose insurance, reinsurance, banking, investment and digital asset sectors use AI in very different ways and at very different levels of materiality.

Our response concentrates on the point where guidance meets software. Regulated entities will rarely fail to agree with the eight principles. The harder question is whether, on any given day, they can show that their AI controls are operating as designed. The Authority's fifth outcome, adequate evidence to enable internal assurance and supervisory review, is where that question is answered, and it is where we believe this Guidance Note can add the most practical value.

OUR CORE PROPOSITION

Evidence by Default. In a well-engineered AI system, every control produces its own evidence as a by-product of running. Ingestion gateways log what they redacted. Verification layers record what passed, what failed and where failures were routed. Approvals are captured inside the deployment workflow rather than in inboxes. Supervisory review then becomes a query against records that already exist, rather than a project to reconstruct them.

Summary of recommendations

1. **Recognise system-generated evidence explicitly.** State that contemporaneous, system-generated control records are a preferred form of evidence for material and higher-impact AI, while confirming that documentary evidence remains acceptable for lower-impact use cases.
2. **Publish an illustrative minimum evidence set by materiality tier.** A short, non-binding annex showing what "adequate" evidence looks like for lower-impact, material and higher-impact use cases would make the proportionality principle easier to apply consistently across sectors.
3. **Specify indicative minimum fields for AI inventories.** Common fields (owners, purpose, materiality rating, data sensitivity, third-party dependencies, oversight model, last validation date) would improve comparability without prescribing a system.
4. **Strengthen agentic AI expectations with enforceable boundaries.** Building on the Authority's statement that tool and API controls should not rely solely on prompts, we propose action tiering, permission manifests, distinct non-human identities, tested suspension controls and delegation lineage (Section 3).
5. **Treat material vendor model updates as change events.** Where a third-party model update can alter outputs without any internal code change, re-verification proportionate to materiality should be expected, supported by contractual notice of model changes.

6. **Connect AI incident records to PIPA obligations.** Decision lineage should be sufficient to support a timely assessment of whether a personal information breach is likely to adversely affect individuals, as the Personal Information Protection Act 2016 (PIPA) requires.
7. **Add worked sector examples.** In answer to the Authority's third question, we suggest short examples for claims triage, underwriting support, AML/ATF transaction monitoring, customer-facing generative assistants and agentic back-office automation.

About Laurenceo

Laurenceo is an independent research initiative based in Hamilton, Bermuda, focused on how AI governance requirements translate into system architecture and day-to-day operation in regulated environments. Its published research framework forms the basis of this response:

- **Building Compliant AI Systems in Bermuda** (15 August 2026): the case for designing compliance into AI systems from the outset in Bermuda's regulatory environment.
- **A Reference Model for AI Governance** (30 August 2026): four structural control layers, namely ingestion controls and privacy boundaries, deterministic output verification, audit trails and decision lineage, and modular isolation.
- **Operating AI Governance in Production** (29 September 2026): four operating disciplines, namely a living AI inventory, evidence by default, drift monitoring, and change control with incident response.

The response is prepared by Laurenceo's founder, Jorai Laurenceo, a Business Systems Analyst whose background spans systems analysis and workflow automation through a placement with Rize Technologies, a Bermuda managed service provider; immersion in reinsurance through a placement with Aeolus Capital Management; and study in technology management and enterprise architecture as a Bachelor of Commerce (Honours) candidate in Technology Management at Ontario Tech University.

Laurenceo submits this response in its own capacity. The views expressed are Laurenceo's and do not represent the views of any employer, placement host or academic institution. Laurenceo has no commercial relationship with any entity regulated by the Authority in connection with this submission.

Response to the Eight Principles and Five Core Outcomes

We support each of the eight principles. For each, we set out our position, the production control that gives it practical effect, and where relevant a targeted enhancement for the Authority's consideration.

Principle	Position	How it is operationalised in production	Suggested enhancement
1. Technology neutrality	Support	Deterministic controls (input validation, output verification, logging) apply to any automated decisioning, whether or not it meets a definition of "AI". This keeps expectations tied to risk and outcome rather than to labels.	Confirm that controls proven for conventional models and rules engines can be reused for AI, reducing duplication and the incentive to relabel systems to avoid scrutiny.
2. Accountability	Strongly support	Every inventory entry names a business owner and a technical owner. Escalation paths for every monitoring threshold resolve to a named person.	Note that an alert with no accountable owner is not a control, and that accountability should be traceable from the board to each material use case.
3. Proportionality	Strongly support	A materiality rating in the inventory sets the depth of validation, monitoring, oversight and evidence for each use case.	Publish an illustrative minimum evidence set per tier (see Section 4) to support consistent application across sectors and entity sizes.
4. Integration with current frameworks	Strongly support	AI risks are recorded in existing risk registers, operational resilience, outsourcing and ICT frameworks rather than in a parallel structure.	Provide a short cross-reference showing where AI evidence maps to existing sectoral codes and returns.
5. Robust lifecycle management	Support	Gated stages from selection to retirement, each producing a signed approval record inside the deployment workflow.	Expressly include decommissioning evidence (data disposal, credential revocation, archived lineage) as part of the lifecycle.
6. Responsible and ethical data use	Strongly support	An ingestion gateway classifies inputs and redacts or blocks personal information before it reaches a model, logging each decision.	Reference PIPA expressly where personal information enters prompts, training data, retrieval stores or agent memory.
7. Explainability and transparency	Support	Decision lineage records inputs, model version, prompt or configuration version, verification result and human action for each material decision.	Recognise lineage as a practical route to explainability where model internals are opaque, particularly for third-party models.
8. Adaptability	Strongly support	Modular isolation lets entities replace a model or vendor without re-engineering core platforms, and re-verify only what changed.	Commit to periodic review of the Guidance Note, informed by incident themes observed across the sector.

The five core outcomes

The five outcomes provide a clear test of whether governance is working. We map each to the evidence a production system can generate automatically.

Outcome	What "working" looks like in production	Evidence generated by default
1. Accountable governance and risk ownership	Every AI use case, including AI embedded in vendor software, is in a living inventory with named owners and a materiality rating. Board reporting draws directly from that inventory.	Inventory change history; owner attestations; board management information generated from the inventory rather than compiled by hand.
2. Proportionate lifecycle management	Control depth follows materiality. Higher-impact systems pass independent validation before release; lower-impact tools operate under an approved-use policy and access controls.	Stage-gate approvals captured in the deployment pipeline; validation reports linked to the exact model and configuration version released.
3. Reliable, explainable and appropriately overseen outcomes	Outputs are verified against schemas and business thresholds before they execute. Failures route to a deterministic fallback or a human queue.	Verification pass and fail records; override rates; drift metrics against validation baselines; human review records.
4. Secure and resilient deployment, including third parties	Prompt injection, data leakage and unsanctioned "shadow AI" are addressed through existing ICT controls. Vendor changes trigger re-verification.	Gateway block logs; red-team and injection test results; vendor change notices linked to re-verification results; suspension test records.
5. Adequate evidence for internal assurance and supervisory review	Any material decision can be reconstructed and any control shown to be operating, on request, without manual reconstruction.	Tamper-evident, hash-chained audit trail; an evidence index mapping each record to the requirement it supports.

Responses to the Authority's consultation questions

Q1. Does the Guidance Note appropriately explain how existing requirements apply to AI?

Largely, yes. The Paper is clear that it creates no new statutory obligations, that accountability is not transferred by outsourcing, and that alternative controls are acceptable where they achieve consistent outcomes. The area that would benefit most from further explanation is evidence: what form of record the Authority would regard as sufficient for a material use case, and how long such records should be retained. Clarity here would reduce the risk that entities default to after-the-fact documentation, which is slower to produce and less reliable than contemporaneous system records.

Q2. Is the materiality and risk assessment approach workable across sectors?

Yes, and the distinction between materiality and risk is particularly valuable. The Paper rightly notes that a lower-materiality use case may still present higher risk, for example where it involves sensitive information. We suggest the assessment be recorded as two separate ratings in the inventory (materiality of the decision, and inherent risk of the data, autonomy and access involved), with the higher of the two setting control depth. This keeps the approach workable for smaller entities while preventing low-materiality tools with broad data access from receiving too little attention.

Q3. Would further sector-specific examples be useful?

Yes. We recommend brief worked examples for (a) claims triage in insurance, (b) underwriting and pricing support in reinsurance, (c) AML/ATF transaction monitoring, (d) customer-facing generative assistants in banking or investment business, and (e) agentic automation of back-office processes. Each example could show the materiality rating, the oversight model, and the minimum evidence an entity would be expected to hold.

Concrete Enhancements for Higher-Impact and Agentic AI Controls

The Paper's treatment of agentic AI is well judged. It expects proportionate limits on autonomy, objectives, duration, tools, data access and permissions; bounded environments; records of actions, tool calls, access decisions, approvals and exceptions; and monitoring for unusual tool use, scope creep and repeated failure. Most importantly, it states that controls over tools and APIs should not rely solely on prompts or model instructions. The enhancements below are intended to give that last principle concrete, testable form.

E1 ACTION TIERING

Classify every action an agent can take into three tiers: **read** (retrieve information), **reversible write** (draft, stage or queue), and **irreversible or external effect** (payments, trades, policy binding, claims settlement, customer communications, changes to records of account). Irreversible actions in higher-impact use cases should require either pre-execution human approval or a deterministic verification gate, never model judgement alone.

E2 PERMISSION MANIFESTS ENFORCED OUTSIDE THE MODEL

Each agent should have a declared, versioned manifest of the tools, data sources and actions it may use, held in the AI inventory and enforced by the platform at the point of execution. Any change to the manifest is a change event subject to approval. A model that is instructed not to use a tool, but technically can, is not bounded.

E3 DISTINCT NON-HUMAN IDENTITIES AND LEAST PRIVILEGE

Agents should act under their own scoped credentials, never under a human user's session or a shared service account. This makes every action attributable, allows access to be revoked instantly, and lets existing identity and access reviews cover agents without new processes.

E4 TRANSACTION, VELOCITY AND DURATION LIMITS

Hard limits on value per action, number of actions per period, cumulative exposure and run time, enforced deterministically. Breaching a limit should halt the agent and escalate, not merely log a warning.

E5 TESTED SUSPENSION CONTROLS

The ability to suspend an agent or automated decision process should be tested on a defined schedule proportionate to materiality, with a defined safe state (for example, revert to a deterministic rule set or manual processing). The test itself should generate an evidence record. A suspension capability that has never been exercised cannot be relied upon.

E6 DELEGATION LINEAGE FOR MULTI-AGENT SYSTEMS

Where agents call other agents or external tools, the audit trail should preserve the full chain: which agent initiated a task, under whose authority, with which inputs, and which downstream calls followed. Without this, accountability dissolves at the first handoff.

E7 MEMORY AND RETRIEVAL GOVERNANCE

Agent memory, retrieval stores and conversation logs are data stores and should be treated as such: classified, subject to retention rules, covered by PIPA where they contain personal information, and capable of targeted deletion.

E8 DETERMINISTIC OUTPUT VERIFICATION FOR HIGHER-IMPACT GENERATIVE AI

Before a generative output executes or reaches a customer, it should pass schema validation, business-rule and threshold checks, and, where it cites facts or policy terms, a check against the authoritative source. Failures route to a defined fallback. This addresses hallucination risk with a control whose operation can be evidenced, rather than relying on model quality alone.

E9 VENDOR MODEL UPDATES AS CHANGE EVENTS

Third-party model updates can change outputs without any change to an entity's own code. For material use cases, entities should obtain contractual notice of model changes where available, pin model versions where possible, and re-run a proportionate verification suite on each update. Where vendor transparency is limited, independent output testing is a practical compensating control.

PROPOSED ADDITION TO THE GUIDANCE NOTE

"For agentic AI used in material or higher-impact activities, the Authority expects permissions, limits and suspension controls to be enforced by the systems within which the agent operates, independently of the agent's own instructions, and expects the regulated entity to be able to demonstrate that these controls have been tested."

Operationalising "Adequate Supervisory Evidence"

The Laurenceo Production Control Specification

The Paper lists the information the Authority may request: governance documents, inventories, materiality and risk assessments, validation and testing results, monitoring and drift reports, oversight arrangements, incident records, third-party due diligence and audit findings, among others. The specification below shows how each category can be produced continuously by the system itself. It is offered as an illustrative reference, consistent with the Paper's position that entities may adopt alternative controls that achieve the same outcomes.

4.1 Materiality tiers

Tier	Typical use cases	Control and evidence depth
Tier 1 Lower impact	Internal drafting, summarisation of non-sensitive material, coding assistance, internal search.	Inventory entry, approved-use policy, access controls, ingestion gateway for sensitive data, periodic attestation.
Tier 2 Material	Decision support where a human makes the final call: claims triage, underwriting support, AML/ATF alert prioritisation, customer communication drafts.	All Tier 1 controls, plus pre-deployment validation, output verification, drift monitoring with named escalation, override tracking and decision lineage.
Tier 3 Higher impact	Binding or customer-affecting decisions, pricing, valuation, trading, and agentic systems with write access or external effect.	All Tier 2 controls, plus independent validation, enhanced human oversight, the agentic controls in Section 3, tested suspension and board-level reporting.

The higher of the materiality rating and the inherent risk rating (data sensitivity, autonomy, breadth of access) determines the tier.

4.2 Control specification

ID	Control	Mechanism	Evidence produced by default	Outcome / Tier
LPC-01	Living AI inventory	Register of all AI use cases, including AI embedded in vendor software, with owners, purpose, materiality, data sensitivity, dependencies and oversight model. Discovery checks flag unregistered AI services.	Versioned inventory with change history; unregistered-service alerts and their resolution.	O1 O5 All tiers
LPC-02	Ingestion controls and PIPA redaction gateway	All inputs pass a classification and redaction layer before reaching a model; prohibited data is blocked; sanctioned tools only.	Per-request log of classification, redaction and block decisions, without storing the redacted content.	O4 All tiers

ID	Control	Mechanism	Evidence produced by default	Outcome / Tier
LPC-03	Deterministic output verification	Outputs validated against schemas, business rules and thresholds before execution; schema fallback to a rule set or human queue on failure.	Pass, fail and fallback records linked to each decision; failure-rate trends.	O3 Tier 2 and 3
LPC-04	Tamper-evident audit trail and decision lineage	Append-only, hash-chained records of inputs, model and configuration version, verification result, human action and outcome.	Reconstructable decision history; chain integrity checks run automatically.	O3 O5 Tier 2 and 3
LPC-05	Risk-tiered lifecycle gates	Selection, validation, release and retirement approvals captured inside the deployment workflow, scaled to tier.	Signed approvals bound to the exact version released; validation reports.	O2 All tiers
LPC-06	Continuous control monitoring	Live behaviour compared with validation baselines: accuracy, input drift, override rates, verification failures.	Monitoring dashboards and threshold breach records with owner acknowledgement.	O3 Tier 2 and 3
LPC-07	Change control	Any model version, prompt, manifest or vendor change triggers re-verification proportionate to tier before production.	Change records linked to re-verification results.	O2 O4 All tiers
LPC-08	Agentic boundary enforcement	Permission manifests, action tiering, scoped identities and hard limits (Section 3).	Tool-call and access-decision logs; limit breaches; manifest change history.	O3 O4 Tier 3
LPC-09	Incident response and suspension	Documented fallback; tested suspension; incident records linked to lineage to support PIPA breach assessment.	Suspension test records; incident timeline generated from the audit trail.	O4 O5 Tier 2 and 3
LPC-10	Board and senior management information	Reporting generated from LPC-01 to LPC-09: material use cases, risk ratings, drift, incidents and near misses, third-party dependencies, significant changes.	Periodic reports with a traceable link from each figure to its source records.	O1 All tiers, scaled

4.3 The evidence record

Each control writes a standard evidence record. A common structure allows evidence to be indexed against requirements, so that an internal auditor or supervisor can ask a question and receive an answer drawn from records that already exist. An illustrative record:

```
{
  "record_id": "evd-2026-10-000184",
  "control_id": "LPC-03",
  "use_case_id": "UC-CLM-007 (claims triage, Tier 2)",
  "timestamp_utc": "2026-10-14T13:42:08Z",
  "model_version": "vendor-model 2026-09 (pinned)",
  "config_version": "prompt v14 / manifest v3",
  "event": "output_verification",
  "result": "fail: reserve estimate above authority threshold",
  "action": "routed to human review queue",
  "owner": "Head of Claims (business) / Platform Lead (technical)",
  "maps_to": ["Outcome 3", "Principle 7", "BMA evidence: monitoring, oversight"],
  "prev_hash": "9flc...e2a7",
  "record_hash": "4b0d...71c3"
}
```

4.4 Why evidence by default matters to the Authority

Reliability

Contemporaneous system records are harder to alter and less prone to omission than evidence assembled after a request is received.

Proportionality

Smaller entities using lower-impact tools need only inventory and access records, while the same structure scales to higher-impact systems without redesign.

Supervisory efficiency

Reviews can focus on whether controls are effective, not on whether documents can be found.

Faster incident response

Decision lineage shortens the time needed to establish scope and impact, including under PIPA breach notification requirements.

Formal Closing and Submission Metadata

Laurenceo thanks the Authority for the opportunity to comment and commends the consultative and proportionate approach taken in the Paper. The Guidance Note has the potential to give Bermuda one of the clearest AI supervisory frameworks of any international financial centre, precisely because it builds on existing obligations rather than adding a parallel regime.

Our central recommendation is that the Authority make explicit what good evidence looks like. If the Guidance Note signals that system-generated, contemporaneous evidence is the expected standard for material AI, and illustrates the minimum evidence set by tier, regulated entities will build controls that prove themselves continuously rather than at audit time.

Laurenceo would welcome the opportunity to discuss this response, to contribute to sector-specific examples, or to participate in any follow-up industry engagement the Authority may convene.



Jorai Laurenceo

Founder, Laurenceo | Business Systems Analyst

Hamilton, Bermuda | laurenceo.com | jorai@jorailaurenceo.com

Submission metadata

Consultation	Consultation Paper: Guidance Note, The Responsible Use of Artificial Intelligence in Bermuda's Financial Services Sector
Issuing authority	Bermuda Monetary Authority
Date issued	14 August 2026
Feedback deadline	30 October 2026
Submitted to	policy@bma.bm
Respondent	Laurenceo, an independent AI governance research initiative, Hamilton, Bermuda
Respondent category	Other stakeholder (independent research)
Contact	Jorai Laurenceo, Founder jorai@jorailaurenceo.com
Document reference	LAU-BMA-AI-2026-01
Confidentiality	Not confidential. Laurenceo consents to the Authority publishing or summarising this response with attribution. Submitted in Laurenceo's own capacity; not legal, regulatory or financial advice.
Supporting research	laurenceo.com/building-compliant-ai-001 laurenceo.com/a-reference-model-002 laurenceo.com/operating-ai-governance-003