

DABUDA LIMITED

AI ASSURANCE EVIDENCE SERIES

Enterprise AI Release Governance Case Study

Illustrative control model for AI release readiness, stakeholder decision records, and post-release monitoring in a regulated enterprise.

ILLUSTRATIVE CASE STUDY - NOT A NAMED CLIENT ENGAGEMENT

Version: April 2026

Tone: Executive, evidence-led, procurement-ready

Prepared for controlled review. Not a certification, endorsement, or named-client reference.

ILLUSTRATIVE CASE STUDY

AI Release Governance for a Regulated Enterprise

Document status

This case study is illustrative. It shows how DaBuDa would structure release governance evidence for a regulated enterprise. It does not represent a named client engagement.

The Problem

A regulated enterprise intends to deploy an AI system to support customer interactions, decision-support workflows, and operational efficiency. The organisation has existing software-release governance, but AI introduces model behaviour, data drift, explainability, and fairness risks that are not fully covered by traditional QA gates.

Challenge	Governance implication
Lack of structured AI release governance	Decision-makers cannot consistently determine whether the system is ready for controlled release.
Limited visibility into model behaviour	Failure modes, edge cases, and inconsistent outputs may remain hidden.
No defined release readiness criteria	Approval may rely on subjective confidence rather than measured evidence.
Regulatory exposure from unvalidated outputs	The organisation may be unable to evidence appropriate controls, review, or risk acceptance.

Release Readiness Criteria

DaBuDa defines release gates that link technical evidence to accountable business decisions. Criteria are set before testing begins and are adjusted only through documented governance approval.

Gate	Evidence required	Example release threshold
Functional accuracy	Scenario pass rate and expert review.	Meets agreed accuracy target for each intended use category.
Consistency	Repeat testing across semantically equivalent prompts.	No material variation in high-impact outputs.
Bias and fairness	Subgroup scenarios, equality review, and disparity analysis.	No unmitigated disproportionate impact in tested scenarios.
Explainability	Human-readable explanations and traceability samples.	Outputs can be reviewed and challenged by accountable stakeholders.
Auditability	Logs, model version, prompt/version history, decision records.	Evidence sufficient for internal audit and change-control review.

Scenario Testing and Control Evidence

Testing is designed to expose behaviour under normal conditions, high-volume use, model updates, edge cases, malicious inputs, and integration failures.

Testing activity	Evidence produced
High-volume scenario simulation	Performance, latency, throughput, and failure-rate metrics.
Edge-case validation	Known limitation register and failure-mode analysis.

Testing activity	Evidence produced
Regression testing across model updates	Version comparison, before/after output analysis, and reproducibility logs.
Adversarial input testing	Prompt injection attempts, data leakage checks, misuse resistance, and safe refusal behaviour.

Stakeholder Decision Records

AI release decisions should be documented in a way that allows business owners, risk owners, technical teams, and auditors to understand what was approved, why it was approved, and what residual risk remains.

Record	Minimum content
Release decision note	System version, intended use, scope, release type, evidence reviewed, release outcome.
Risk acceptance statement	Residual risks, accountable owner, mitigation plan, review date, and approval status.
Approval checkpoint	Named decision-makers, role, date, condition of approval, and dependencies.
Cross-functional sign-off	Product, engineering, security, legal, data protection, operations, and risk review.

Post-Release Monitoring Alignment

AI assurance continues after deployment. Monitoring should detect changes in output quality, user behaviour, data patterns, model behaviour, and operational impact.

Monitoring area	Purpose
Performance monitoring	Confirm that the system continues to meet release criteria in production.
Drift detection	Identify when data or model behaviour diverges from tested conditions.
Incident tracking	Record harmful outputs, near misses, escalation failures, or user complaints.
Periodic re-evaluation	Re-run assurance after model updates, policy changes, new use cases, or material incidents.

Example Governance Outcome

Status: Conditional Enterprise Release

- Release permitted only within defined business workflows and approved user groups.
- AI outputs remain advisory and must not constitute final regulated decisions without human review.
- Release owner must maintain evidence pack, issue register, monitoring dashboard, and version history.
- Any model, prompt, data, or integration change must trigger regression testing and governance review.
- Full release deferred until escalation accuracy, output consistency, and incident-response playbooks are validated.

Important Note

Illustrative scope disclaimer

This case study is illustrative only. It does not represent a specific client engagement, formal regulatory opinion, procurement reference, or verified deployment outcome.

Guidance Alignment

This document is aligned to current UK public-sector AI assurance expectations and uses the following references as governance context. It does not claim certification, accreditation, formal approval, or legal compliance.

- UK Government: AI Playbook for the UK Government, 10 February 2025.
- UK Government: Algorithmic Transparency Recording Standard guidance for public sector bodies, 8 May 2025.
- UK Government: Data and AI Ethics Framework, updated 18 December 2025.
- Cross-Government Testing Community: AI Testing and Assurance Framework for Public Sector.
- UK Government: Trusted third-party AI assurance roadmap, 3 September 2025.

Note

For regulated enterprises, this pack should be mapped to the organisation's sector-specific obligations, internal risk taxonomy, audit standards, and model-risk management process.