

NOETFIELD TRUST LEDGER – SAMPLE BOARD REPORT

Illustrative template for a Q1 2026 reporting period

Cover / Front Page

Organisation: Example Institution

Committee: Board Risk Committee

Report Title: Noetfield Trust Ledger – Sample Board Report

Reporting Period: Q1 2026

Prepared By: AI Governance Function

Contact: ai-governance@exampleinstitution.com

Version: Sample v1.0

Purpose of this document

This document is an illustrative example of a Noetfield Trust Ledger board report.

It shows how material AI systems, significant changes, incidents, controls and board actions can be summarised in a concise, human-readable format for:

- Board directors and board committees
- Executive management
- Internal audit and risk functions
- Supervisors and regulators (where appropriate)

This sample does not represent a real institution and is provided for demonstration purposes only.

1. EXECUTIVE SUMMARY

1.1 Purpose of the Executive Summary

The Executive Summary gives the Board a one-page view of:

- The most important AI governance messages this quarter (“Headlines”)
- The specific decisions requested from the Board or its committees
- Any change in overall AI risk level that directors should be aware of

It is designed so that a director can understand, in a few minutes, where to focus attention and what is being asked of them.

1.2 Headlines – What the Board Needs to Know

1. Stability of AI credit assistant (S-01).

The AI-driven credit decision assistant for retail customers remains stable.

- Parameter changes introduced last quarter (for example, maximum auto-approved limit increases) have not led to any material deterioration in key risk indicators.
- Delinquency, affordability and fairness metrics remain within the appetite previously approved by the Board Risk Committee.

2. New AI vendor-screening system (S-03) ready for launch.

A new AI-based system for high-risk vendor screening has completed internal testing and governance review.

- The AI Governance Committee recommends a controlled production launch, subject to the safeguards and monitoring thresholds described later in this report.
- Because the system can affect onboarding of critical vendors, the final go-live is submitted to the Board Risk Committee for approval.

3. One medium-severity near miss successfully caught.

A medium-severity near miss (I-01) occurred in the credit assistant (S-01).

- A suggested credit adjustment would have breached internal affordability guidance.

- The issue was caught by a human reviewer before customer impact, demonstrating that the human-in-the-loop control (C-01) is working.
- Prompts and configuration have been updated, and staff received targeted training.

1.3 Board Decisions Requested

For this quarter, the Board (or Board Risk Committee) is requested to:

1. BA-01 – Approve deployment of S-03 (AI vendor screening) into production.
 - Decision type: Approve / Amend / Reject
 - Recommended decision: Approve with conditions, including:
 - Quarterly review of false-positive and false-negative rates
 - Explicit escalation if defined thresholds are breached
 - Confirmation that human reviewers can override AI outcomes where risk is uncertain
2. BA-02 – Confirm decommissioning of S-07 and archival of logs.
 - S-07 is a legacy NLP classifier that has been replaced by S-02 (customer support triage).
 - Decision type: Confirm
 - Recommended decision: Approve, as there is no remaining business dependency and all required logs have been archived in line with retention policies.

1.4 How to Read the Rest of the Report

- Section 2: Which AI systems are considered “material” and therefore in scope
- Section 3: What has changed (new systems, parameter changes, decommissioning)
- Section 4: Incidents and near misses that directors should know about

- Section 5: Key controls and test results, including any gaps
- Section 6: High-level view of provider usage and infrastructure
- Section 7: Detailed breakdown of the Board actions requested
- Section 8: Optional annexes and additional detail (for a real deployment)

2. AI SYSTEM INVENTORY SNAPSHOT

2.1 Purpose of the Inventory

This section lists AI systems that are “material” for governance purposes – that is, systems that:

- Can move money, affect customers, impact markets, or
- Trigger meaningful regulatory, operational or reputational risk

It is not an exhaustive technical inventory, but a governance-scoped subset.

2.2 Definition – “Material AI System”

For the purposes of this report, a “material AI system” is one that meets at least one of the following:

- Supports or executes decisions that affect financial exposure (for example, credit limits)
- Influences regulated processes (for example, KYC/KYB, vendor risk, sanctions screening)
- Directly affects customer outcomes at scale (for example, complaints routing, triage)
- Is likely to attract supervisory or public interest if it fails

2.3 Inventory Table (Illustrative Example)

You can paste the table below directly into Word and format as a normal table.

Table 1 – Material AI systems in scope (Q1 2026)

ID	System / Use Case	Business Owner	Model / Provider	Infrastructure	Data Sensitivity	Risk Tier	Status
S-01	Credit decision assistant – retail limits	CRO	Internal model + hosted frontier API	GPU-backed cloud	High	Tier 2	Production
S-02	Customer support triage – inbound message s	COO	Hosted LLM with internal guardrails	Cloud CPU	Medium	Tier 3	Pilot
S-03	Vendor risk screening – high-risk vendors	CRO	Composite (internal + third-party)	GPU-backed cloud	Medium	Tier 2	Pre-production

2.4 Key Points for the Board

- S-01 remains the most sensitive system due to direct credit exposure and customer impact.
- S-03 (vendor screening) is moving from design/test into a proposed controlled production launch.
- S-02 is still treated as a pilot; lessons learned from incidents and customer feedback will be used before scaling further.

3. SIGNIFICANT CHANGES THIS PERIOD

3.1 Why Changes Matter

Changes to AI systems – including new deployments, parameter changes, and decommissioning – are a common source of emerging risk.

The Board should be aware of what changed, why, and how it was governed.

3.2 Change Log (Illustrative Example)

Table 2 – Significant AI-related changes in Q1 2026

Change Type	System ID	Description of Change	Rationale	Governance Actions	Effective Date
New	S-03	Introduction of AI-based vendor screening for high-risk vendors.	Reduce manual review time; increase consistency and coverage.	Reviewed by AI Governance Committee; recommended for Board Risk Committee approval with conditions on monitoring.	2026-01-15
Parameter change	S-01	Max auto-approve d credit limit increase raised from +10% to +15% for low-risk customers.	Align with peer practice; support growth while within risk appetite.	Impact analysis presented to Risk Committee; endorsed with weekly drift and delinquency monitoring requirements.	2026-01-28

Decommissioned	S-07	Legacy NLP classifier decommissioned and replaced by S-02.	Consolidate into a unified stack with stronger monitoring.	Decommission plan approved; logs archived for 7 years under retention standard AIG-02.	2026-02-03
----------------	------	------------------------------------------------------------	------------------------------------------------------------	----------------------------------------------------------------------------------------	------------

3.3 What the Board Should Take From This

- No change this period moves the institution outside the previously approved risk appetite, but:
 - The parameter change in S-01 should continue to be monitored closely.
 - The introduction of S-03 requires explicit Board-level approval given its impact on vendor onboarding for high-risk vendors.

4. INCIDENTS, EXCEPTIONS, AND NEAR MISSES

4.1 Purpose of this Section

This section highlights AI-related events that are important for governance, including:

- Confirmed incidents
- Policy exceptions
- Near misses where controls worked but risk was present

The goal is transparency: directors should see where AI or automation challenged existing safeguards, even if there was no external impact.

4.2 Incident Table (Illustrative Example)

Table 3 – AI incidents and near misses in Q1 2026

Ref	Date	System ID	Severity	Description	Root Cause	Actions Taken	Status
I-01	2026-01-07	S-01	Medium	Model suggested a credit adjustment breached internal affordability guidance; caught by human reviewer in time.	Prompt design gap and missing explicit guardrail.	Updated prompts and configuration; added affordability constraint; targeted reviewer training; added template to playbook.	Closed

4.3 Regulatory Reporting and Visibility

- Under current regulatory thresholds, no events required external notification this quarter.
- Incident I-01 is classified as a board-visible near miss, demonstrating that:
 - The human-in-the-loop control is functioning, and
 - Certain guardrails still need refinement in upstream configuration and design.

5. CONTROLS & TESTING

5.1 Purpose of this Section

The Board needs assurance that key AI controls:

- Are clearly defined
- Have owners
- Are tested regularly
- Produce understandable results (effective / partially effective / gaps)

This section focuses on governance-relevant controls rather than exhaustive IT controls.

5.2 Frameworks and Standards (Illustrative)

Controls in this report are mapped, at a high level, to:

- Internal AI governance standard AIG-01 (design, deployment, monitoring)
- Internal retention and logging standard AIG-02
- Sector guidance and supervisory expectations relevant to AI and model risk

5.3 Control Table (Illustrative Example)

Table 4 – Key AI governance controls and test results

Control ID	Control Description	Systems in Scope	Owner	Testing Frequency	Last Test Date	Result / Key Finding
C-01	Human-in-the-loop review for all high-impact AI credit decisions.	S-01	Risk	Each batch	2026-02-20	Effective – all high-impact decisions reviewed; near miss I-01 detected and stopped before

						customer impact.
C-02	Pre-deployment configuration and prompt review for Tier 1–2 systems.	S-01, S-03	AI Governance	Per deployment	2026-02-05	Minor findings – documentation strengthened; actions completed within the period.
C-03	Quarterly model performance and drift review with challenge from Risk.	S-01	Model Risk	Quarterly	2026-02-12	Effective – performance within approved tolerance; no recalibration required this quarter.

5.4 Control Gaps and Follow-up Actions

- Documentation of configuration reviews (C-02) has been made more explicit, including:
 - Clear sign-off points
 - Storage of review artefacts alongside the Trust Ledger entry
- Extending quarterly drift review (C-03) to S-03 is planned once S-03 is live and has sufficient data.

6. PROVIDER & INFRASTRUCTURE VIEW

6.1 Why the Board Sees This

Many AI systems now rely on:

- External frontier or hosted AI providers
- GPU-backed cloud infrastructure
- Complex chains of data, models, and APIs

This section provides a board-level summary of how these components are used and monitored.

6.2 Frontier / Hosted AI Services

Key metrics (illustrative for Q1 2026):

- Total API calls: 1.8 million (approximately +12% vs prior quarter)
- High-risk use cases using hosted AI: 2 systems (S-01 and S-03)
- AI-related incidents linked to hosted services: 1 (I-01, resolved as a near miss)

Hosted frontier AI services are used in:

- S-01 (Credit decision assistant)
- S-02 (Customer support triage)
- S-03 (Vendor screening)

Safeguards include:

- Contractual clauses on data use and audit rights (aligned with AIG-03)
- Internal prompt and configuration reviews (C-02)
- Monitoring and logging integrated with broader information security standards

6.3 GPU-backed Workloads

Key metrics (illustrative for Q1 2026):

- GPU hours consumed: 420 hours (+8% vs prior quarter)
- New models deployed using GPU-backed infrastructure: 1 (S-03)
- Governance reviews performed: 3 (covering performance, security, and risk)

GPU-backed workloads run in a restricted cloud environment with:

- Network segmentation
- Identity and access controls
- Centralised logging and monitoring

These measures are aligned with the institution's core cybersecurity framework.

6.4 Executive Access and Device Posture

- Access to AI systems and the Trust Ledger by board members and senior executives is restricted to managed devices.
- Controls include:
 - Enforced encryption
 - Strong authentication (for example, MFA)
 - Central logging of access activity
- Temporary exceptions (for example, access from unmanaged or personal devices) require:
 - Explicit approval
 - Time-bound access
 - Additional monitoring

This is included to reassure the Board that AI-related information and decisions are not exposed via weak endpoints.

7. BOARD ACTIONS & APPROVALS

7.1 Purpose of this Section

This section consolidates all decisions requested from the Board or committees into a single table, so that:

- Directors can see what is being asked,
- On what basis, and
- By when a decision is needed.

7.2 Board Action Table (Illustrative Example)

Table 5 – Board actions for Q1 2026

Item	Description	Request	Recommended Decision	Deadline
BA-01	Approve deployment of S-03 (AI vendor screening) into controlled production with defined thresholds and human-review steps.	Approve / Amend / Reject	Approve with conditions – quarterly review of error rates; escalation triggers; review of vendor-impact metrics.	2026-03-10
BA-02	Confirm decommissioning of S-07 and archival of logs and	Confirm	Approve – no remaining dependency; replacement system S-02	2026-03-10

documentation
for
audit/supervisor
y purposes.

operational;
retention
requirements
satisfied.

7.3 How Decisions Are Recorded

- Formal Board decisions can be captured:
 - In Board or committee minutes, with cross-reference to BA-01 and BA-02
 - Or in a dedicated annex to this Trust Ledger report, where sign-offs and conditions are recorded

For a live Trust Ledger, these references become part of the auditable chain of governance for AI.

8. ANNEX (ILLUSTRATIVE)

8.1 Purpose of the Annex

The annex is optional in this sample, but in a live deployment it can hold:

- Detailed system profiles for Tier 1–2 systems
- Full incident and near-miss reports
- Control-to-regulation mappings
- A glossary for non-technical directors

8.2 Examples of Annex Content

- Annex A – System profiles

For each Tier 1–2 system (for example, S-01 and S-03), provide:

- Scope and purpose
- Data sources and sensitivity
- Key safeguards and limits
- Metrics and thresholds
- Ownership and escalation paths
- Annex B – Incident reports

For entries like I-01:

- Timeline of events
- Detailed root cause analysis
- Lessons learned
- Follow-up actions and their status
- Annex C – Regulatory and standards mapping

Mapping of key controls to:

- Sectoral AI and model risk guidance
- Internal policies and standards
- Any applicable AI-specific regulations
- Annex D – Glossary

Simple explanations of:

- “Material AI system”
- “Risk tier”
- “Near miss”
- “Guardrail”

- “Drift review”
and other terms used in the report.