

Orchestra AI

A · I · T H A T · S U R V I V E S · S C R U T I N Y

AI GOVERNANCE ADVISORY · CASE STUDY 02

AI Governance Advisory — *Findings and Recommendations.*

Prepared for: [Client] · Testing Services & SaaS · Anonymised

This report documents the findings of a two-week AI governance advisory engagement. It identifies five governance gaps, provides specific recommendations against each, records the five-component AI Policy Framework delivered, and provides an implementation roadmap for the recommended actions.

ENGAGEMENT TYPE	DURATION	SECTOR	DELIVERED BY	DATE
AI Governance Advisory	2 weeks	Testing Services · SaaS	Orchestra AI	2026

I.

Executive Summary

This engagement was initiated following enterprise pilot interactions in which the client's AI-powered compliance checking product encountered governance questions it could not answer. The product's AI capability was not in question. The absence of a documented AI governance framework prevented the product from proceeding through enterprise procurement.

A two-week advisory engagement was conducted covering discovery of the current AI governance state, gap identification, and delivery of an AI Policy Framework addressing all identified gaps.

GAPS IDENTIFIED

5

3 Critical · 2 High

ENGAGEMENT DURATION

2 wks

Discovery + Framework Design

FRAMEWORK COMPONENTS DELIVERED

5

All gaps addressed

Overall assessment

The client's AI product demonstrates strong technical capability and sound product design intent — particularly the human-in-the-loop (HITL) design placing validation responsibility with the customer. However, the governance infrastructure required to make that HITL meaningful — documented criteria, explainability standards, defined AI scope, and policy-anchored traceability — was absent across all five dimensions assessed.

The AI governance posture at the start of engagement was **undocumented across all five dimensions**. The product was not enterprise-deployable in its pre-engagement state. At the close of the engagement, all five gaps have been addressed through the delivered AI Policy Framework.

Central finding: AI governance cannot be an afterthought for scaling. The enterprise pilot failures were not a product problem — they were a governance infrastructure problem. The framework delivered in this engagement closes that gap and enables enterprise onboarding.

Priority actions

The following three actions are recommended for immediate implementation before the next enterprise pilot engagement:

- 1.** Publish the AI Governance Policy (Recommendation R-01) — this is the document enterprise procurement asks for first.
- 2.** Integrate the Customer Review Documentation Package (Recommendation R-02) into the product output — every AI finding must be accompanied by the review documentation before the next pilot.
- 3.** Document AI Generation Controls and Scope (Recommendation R-03) — enables enterprise clients to see exactly what AI does in their environment.

II.

Governance Assessment *Findings.*

Five governance gaps were identified through the discovery process. Each gap maps directly to a question raised in enterprise pilot interactions that the product could not answer. Severity is assessed against the Navigator Framework risk criteria and enterprise procurement standards for AI products operating in compliance contexts.

F-01 · CRITICAL · GOVERNANCE**No documented AI governance framework**

FINDING

The AI product operates under no documented governance policy. No defined rules exist for what the AI is permitted to generate, how its outputs are supervised, or who holds accountability for those outputs. This is the foundational governance requirement for any AI product entering enterprise procurement — and it was entirely absent.

EVIDENCE

Unable to produce a governance policy document when requested by enterprise procurement during pilot. No internal documentation of AI operating boundaries, accountability structure, or oversight design.

IMPACT

Enterprise clients with their own AI risk governance obligations will not onboard a product operating without a documented governance framework. This gap is the primary reason enterprise pilots stalled at procurement stage.

F-02 · CRITICAL · CUSTOMER OVERSIGHT**Customer HITL without documentation to review against**

FINDING

The product's design places the customer as the human-in-the-loop reviewer — a sound governance design. However, customers receive AI-generated findings with no documentation of the criteria, scope, or logic that produced them. Without this documentation, customers cannot perform meaningful review. The HITL is present; the tools to exercise it are not.

EVIDENCE

No customer review documentation package exists. Enterprise pilot contacts asked what documentation their compliance team would receive — no answer was available.

IMPACT

Customer approval of AI findings without documentation constitutes liability transfer without governance. Enterprise legal and risk teams identified this as an unacceptable risk posture in the compliance context.

II. (continued)

Governance Assessment Findings (*continued*).

F-03 • CRITICAL • AI CONTROLS

No documented controls on AI generation scope or criteria

FINDING

No specification documents what the AI checks, how it checks, what criteria it applies, or within what defined scope it operates. The AI functions as an agentic engine — capable of generating compliance findings — but without documented boundaries or criteria governing its operation.

EVIDENCE

No specification documents for AI check logic. No defined criteria document for what constitutes a finding under each compliance framework checked. Check logic was generated by coding agents from informal requirements without formal specification review.

IMPACT

Enterprise clients cannot deploy an AI product in their compliance environment without knowing what that AI does and within what boundaries. Ungoverned AI generating compliance findings inside an enterprise environment creates unacceptable and undocumented risk exposure for the client.

F-04 • HIGH • TRACEABILITY

Traceability controls present but not effective — no policy benchmark

FINDING

The product has traceability controls recording AI process steps. These controls are not effective because there is no documented policy defining what those steps should achieve. Controls without a policy benchmark are process logs, not governance evidence.

EVIDENCE

Traceability control infrastructure confirmed present. Unable to produce the documented policy or standard the controls are designed to enforce when queried in enterprise due diligence.

IMPACT

Existing traceability controls cannot be used as governance evidence in audit or due diligence until a policy benchmark is established. Investment in traceability infrastructure is not yielding its governance value.

F-05 • HIGH • DEVELOPMENT PROCESS

No responsible AI policy for the development process

FINDING

The product was built using coding agents without a spec-driven development process. No policy governed how AI was used in development — no requirements before coding agents began, no defined acceptance criteria, no formal review gates. Enterprise due diligence increasingly covers how an AI product was built, not just how it operates.

EVIDENCE

No specification documents for AI-generated features. Check logic cannot be traced to a reviewed and approved requirement. Test suite generated by the same coding agent as the code it validates — Test Independence Ratio: 0%.

Unable to demonstrate responsible AI development practice to enterprise due diligence. Exposes the organisation to the question "how was this built?" without a structured answer. Ongoing product development without this policy will compound the gap with each new feature.

III.

Recommendations.

Five recommendations correspond directly to the five findings. Each recommendation is stated as a specific action, with implementation guidance and a priority designation. Recommendations R-01, R-02, and R-03 are designated Immediate — required before the next enterprise pilot engagement.

R-01 Publish the AI Governance Policy — *immediately*.

ADDRESSES: F-01 PRIORITY: **IMMEDIATE** OWNER: PRODUCT LEAD + LEGAL

Publish a documented AI Governance Policy covering: what AI is permitted to generate, the oversight structure for AI outputs, named accountability for AI outputs at each stage, how the policy is maintained as the product develops, and how clients are notified of material changes to AI governance.

- Use the AI Governance Policy component delivered in the AI Policy Framework as the base document. Review, approve at leadership level, and publish to the client-accessible documentation before the next enterprise pilot.
- Include a version number and review date. Commit to an annual review cadence at minimum — quarterly recommended for the first year given active product development.
- Make the policy available in the product's enterprise documentation pack and able to be shared with enterprise procurement on request without internal approval delay.

R-02 Integrate the Customer Review Documentation Package *into product output*.

ADDRESSES: F-02 PRIORITY: **IMMEDIATE** OWNER: PRODUCT LEAD + ENGINEERING

Every AI finding delivered to a customer reviewer must be accompanied by the Customer Review Documentation Package: the criteria applied, the scope of the check, the basis of the finding, and the specific elements the reviewer needs to assess. This is not a separate document — it is integrated into the product output for each finding.

- Use the Customer Review Documentation Package templates delivered in the AI Policy Framework. Implement as a structured companion to each AI finding output within the product interface.
- Include a reviewer attestation mechanism — a named reviewer sign-off field that records who reviewed the finding, when, and their assessment. This creates the HITL evidence trail that enterprise governance requires.
- Test with an internal reviewer before the next enterprise pilot. The reviewer should be able to evaluate the finding meaningfully using only the documentation package — without needing to consult the product team.

III. (continued)

Recommendations (continued).

R-03 Document AI generation controls and scope *per check type*.

ADDRESSES: F-03 PRIORITY: IMMEDIATE OWNER: PRODUCT LEAD + ENGINEERING

For every compliance check the AI performs, document: what it checks, the criteria applied, the source data or regulatory reference used, the scope of the check (what it covers and what it does not), and the conditions under which a finding is generated. This becomes the AI Generation Controls document referenced by the policy and provided to enterprise clients in due diligence.

- Use the AI Generation Controls and Scope Definition template from the AI Policy Framework. Complete for all current check types before the next enterprise pilot.
- For each new check type added to the product, complete the controls documentation before the check is made available to clients. This is the governance gate for new AI capability — no client access until documented.
- Review and update this documentation when the underlying AI model is updated or retrained. Model changes may affect what the AI detects and how — the controls document must reflect the current model's behaviour.

R-04 Establish the policy benchmark for traceability controls *within 30 days*.

ADDRESSES: F-04 PRIORITY: 30 DAYS OWNER: ENGINEERING LEAD

Define the documented policy that the existing traceability controls enforce. The controls record what happened — the policy defines what should happen and the standard against which the trace is measured. Once the policy benchmark is established, the existing traceability infrastructure becomes governance evidence rather than a process log.

- Use the Policy Benchmark document from the AI Policy Framework. Review with Engineering Lead to confirm it reflects the actual behaviour of the traceability controls in place.
- Link the policy benchmark to the AI Governance Policy (R-01) so that the traceability controls are clearly referenced as governance evidence in the overarching framework.

R-05 Implement the Responsible AI Development Policy *for all future features*.

ADDRESSES: F-05 PRIORITY: 60 DAYS OWNER: ENGINEERING LEAD + PRODUCT LEAD

Implement a spec-driven development policy for all AI-generated features going forward. No AI-built feature should reach the product without a reviewed and approved specification preceding the coding agent, a human review gate before client deployment, and a traceability record from requirement to code. This applies to all new check types, feature additions, and model updates.

- Adopt the Responsible AI Development Policy from the AI Policy Framework. Communicate to engineering team before the next sprint cycle involving AI-built features.

- Introduce a pre-coding specification review step: before any coding agent begins work on a new compliance check, a human-reviewed specification defining the expected behaviour must be approved. The specification is the governance gate for the coding agent.
- Address the Test Independence Ratio: tests for AI-generated check logic should be produced by a reviewer independent of the coding agent that generated the code. For existing check types, remediate test coverage as part of the 60-day implementation window.

IV.

AI Policy Framework — *Deliverables.*

The following five components were produced and delivered at the close of the engagement. Each component is a ready-to-implement document addressing one or more findings. The framework is designed to be immediately usable — not aspirational policy, but operational governance that can be applied from the day of delivery.

D-01 AI Governance Policy

Documents how AI operates in the product — what it generates, how it is supervised, who is accountable for its outputs, and how the policy is maintained. Formatted for direct publication to enterprise documentation. Addresses F-01. Supports R-01.

DELIVERED · READY TO PUBLISH

D-02 Customer Review Documentation Package

Template set for the documentation accompanying each AI finding to the customer reviewer — criteria applied, scope, basis of finding, and reviewer guidance. Includes reviewer attestation record format. Addresses F-02. Supports R-02.

DELIVERED · READY FOR PRODUCT INTEGRATION

D-03 AI Generation Controls and Scope Definition

Structured template for documenting criteria, scope, and boundaries per check type. Partially completed for current check types during the engagement. Remaining check types to be completed by product team using the template and guidance provided. Addresses F-03. Supports R-03.

DELIVERED · COMPLETION IN PROGRESS

D-04 Policy Benchmark for Traceability Controls

The documented policy standard that the existing traceability controls measure against. Transforms the current process log into governance evidence. Addresses F-04. Supports R-04.

DELIVERED · READY TO IMPLEMENT

D-05 Responsible AI Development Policy

Governs AI use in the development process — specification requirements before coding agents begin, review gates before client deployment, traceability requirements from requirement to code, and Test Independence Ratio requirements. Addresses F-05. Supports R-05.

DELIVERED · READY FOR ENGINEERING ADOPTION

V.

Implementation *Roadmap.*

PHASE	ACTION	TIMELINE	OWNER
Immediate	Publish AI Governance Policy (D-01) to enterprise documentation pack	Before next pilot	Product Lead
Immediate		Before next pilot	

	(D-02) into product output for all client-facing AI findings		Product
Immediate	Complete AI Generation Controls documentation (D-03) for all current check types	Before next pilot	Product Lead
30 days	Establish policy benchmark for traceability controls (D-04). Link to AI Governance Policy.	Within 30 days	Engineering Lead
60 days	Implement Responsible AI Development Policy (D-05) for all future AI-built features. Introduce specification review gate and Test Independence Ratio requirements.	Within 60 days	Engineering Lead + Product
Ongoing	Quarterly review of AI Governance Policy. Update AI Generation Controls when model is updated. Annual third-party governance assessment recommended.	Quarterly	Product Lead

AI governance cannot be an afterthought *for scaling*.

This engagement has delivered the governance foundation. The framework is in place. What follows is implementation — and the commercial advantage that comes from being the AI product in your category that can answer every governance question the enterprise pilot room asks.

RECOMMENDED NEXT STEPS

- 01 Implement the three immediate recommendations before the next enterprise pilot engagement

- 02 Share the AI Governance Policy with the enterprise prospects whose pilots stalled — governance documentation may restart conversations

- 03 Complete the AI Generation Controls documentation (D-03) with the product team using the delivered template

- 04 Consider a follow-up Navigator Audit in 6–9 months to assess implementation and identify any new governance gaps as the product scales

Orchestra AI
rajesh@orquestraai.io
orquestraai.io