orquestraai.io
Version 1.0

Orquestra AI

A · I THAT SURVIVES SCRUTINY

NAVIGATOR AUDIT · CASE STUDY 01

Navigator Audit — *Findings and Recommendations.*

Prepared for: [Client] · IT Services & SaaS · Chennai · Anonymised

This report documents the findings of a four-week Navigator Framework™ audit engagement. It identifies five governance gaps assessed against the Navigator Framework™ Sub-framework 02, provides six specific recommendations, records all deliverables, and provides an implementation roadmap.

ENGAGEMENT	FRAMEWORK	DURATION	DELIVERED BY	DATE
Navigator Audit	Sub-framework 02	4 weeks	Orquestra AI	2026

I.

Executive *Summary.*

This engagement assessed the AI governance posture of a mid-size SaaS product and IT services company building AI agents into its platform and developing a standalone AI agent capability. The company had an active R&D team working with AI models, a product roadmap for AI integration, and a growing base of EU and UK clients — all of whom brought live AI Act transparency obligations.

A four-week Navigator Framework™ Sub-framework 02 assessment was conducted across two teams — the AI R&D team and the Product/SDLC team — evaluating five phase gates and five failure modes. The assessment identified five governance gaps and delivered six recommendations with implementation support across weeks three and four.

GAPS IDENTIFIED	ENGAGEMENT	RECOMMENDATIONS	HITL COVERAGE
5	4 wks	6	100%
2 Critical · 2 High · 1 Medium	Assessment + Implementation	All gaps addressed	0% → 100% in 4 weeks

Overall assessment

Both teams demonstrated strong technical capability in their respective domains. The AI R&D team was building agents at pace. The Product team was managing a live multi-client roadmap. Neither team was failing at its own job. The governance failure was structural: the two teams had no shared process, no integration mechanism, and no audit infrastructure connecting AI development to the product, the delivery lifecycle, or client regulatory obligations.

The governance posture at the start of engagement scored **6/20** against the Navigator Framework™ — At Risk. At the close of the four-week engagement, all five gaps had been addressed through the six delivered recommendations.

> **Central finding:** Every gap identified was a governance, process, or communication failure — not a capability failure. The teams were skilled. The problems were structural. All five were addressable within the engagement window.

Priority actions

1. Activate the R&D-to-Product integration workflow (R-01) — the foundational fix connecting the two teams.

2. Implement HITL approval framework (R-04) — named accountability for every AI output before client-facing deployment.

3. Begin EU/UK explainability documentation (R-03) — live regulatory obligations require immediate action.

II.

Governance Assessment *Findings.*

Five gaps were identified through the four-week assessment across both teams. Each gap maps to one or more Navigator Framework™ failure modes: Accountability, Compliance, Competency, Quality Drift, and Operational Resilience. Severity is assessed against Navigator Framework™ criteria and EU AI Act obligations applicable to the client's active UK/EU client relationships.

F-01 · CRITICAL · ACCOUNTABILITY

No audit trail of any kind across either team

FINDING

No record existed of what AI agents had been built, what decisions they made, on what data, under what constraints, or who had approved them. Zero audit infrastructure existed across both teams. No named human could defend any AI decision — because no record of any AI decision existed.

EVIDENCE

No agent build registry. No approval records for any AI output entering the product. No data lineage documentation. No record of what agents were in development vs. deployed.

IMPACT

Complete inability to respond to any regulatory, client, or board inquiry about what AI the organisation was running, what it had decided, and who had approved it. The most fundamental accountability requirement — that a named human can explain an AI decision — could not be met.

F-02 · CRITICAL · COMPLIANCE

Complete structural silo between AI R&D and Product teams

FINDING

The AI R&D team was building agents the Product team did not know existed, for use cases the Product team had not defined, with no record that any of it had happened. The Product team was managing a live roadmap with no visibility into AI capabilities being built two teams away. No integration process, no shared backlog, no handoff mechanism existed.

EVIDENCE

Product team unable to name what AI agents were in development. R&D team unable to map their agents to live product use cases. No shared documentation, no joint review process, no integration checkpoint in the SDLC.

IMPACT

AI capability being built with no product alignment. AI entering the SDLC without defined requirements, use case validation, or integration governance. Risk of AI features deployed to clients that have not been reviewed for product fit, safety, or regulatory compliance.

II. (continued)

Governance Assessment Findings *(continued).*

F-03 · HIGH · COMPLIANCE

No explainability documentation for EU and UK clients

FINDING

EU AI Act transparency obligations for limited-risk AI systems were in effect for the company's EU and UK client base — including disclosure requirements, accountability documentation, and explainability for automated outputs. No such documentation existed for any AI agent in development or deployment.

EVIDENCE

No explainability templates. No disclosure notices for EU/UK clients using AI-assisted features. No documentation of what data AI agents processed or what decisions they influenced for any client account.

IMPACT

Live regulatory non-compliance for EU and UK client accounts. Any client exercising their right to explanation of an AI-assisted decision would trigger a compliance gap that could not be remediated from existing documentation. Regulatory exposure was immediate, not prospective.

F-04 · HIGH · ACCOUNTABILITY

No human-in-the-loop approval for any AI decision

FINDING

AI agents were being developed and deployed without any defined HITL approval points. No named human was required to review, approve, or attest to any AI output before it entered the product or was demonstrated to clients. HITL coverage at assessment: 0%.

EVIDENCE

No approval records for any deployed AI feature. No defined sign-off process in the SDLC for AI outputs. AI-generated outputs entering client-visible product surfaces without review.

IMPACT

No named accountability for any client-facing AI decision. In the event of an AI error affecting a client, no human approval record exists to establish that appropriate oversight was applied. This is the core governance requirement the Navigator Framework™ is designed to close.

F-05 · MEDIUM · COMPETENCY

Capability mismatch in both directions

FINDING

The Product team lacked AI vocabulary to evaluate what R&D was building or set meaningful requirements for AI features. The R&D team lacked product knowledge to prioritise use cases or understand client context. Neither team could effectively brief the other. The mismatch was creating wasted R&D effort and product risk.

EVIDENCE

Product team unable to specify AI requirements or evaluate AI outputs. R&D team building features without client-validated use cases. No shared vocabulary for AI capability, risk, or evaluation criteria.

IMPACT

Ongoing risk of AI development misaligned to product needs. Governance gaps compounded by communication gaps — even well-intentioned oversight cannot function across a competency mismatch of this magnitude.

III.

Recommendations.

Six recommendations address the five findings. R-01, R-03, and R-04 are designated Immediate — required before any further AI feature enters the product. R-02, R-05, and R-06 are designated within the implementation window of weeks three and four.

R-01 Establish the R&D-to-Product *integration workflow.*

ADDRESSES: F-02 PRIORITY: IMMEDIATE OWNER: PRODUCT LEAD + R&D LEAD

A defined integration workflow connecting AI agent development to the product roadmap — a shared backlog, a joint use-case review cadence, and a handoff checkpoint before any AI output enters the SDLC. No AI feature proceeds to product integration without a product-reviewed use case specification.

→ Establish a weekly joint use-case review between R&D Lead and Product Lead. All AI capabilities under development reviewed against live product requirements before proceeding.

→ Create a shared AI feature backlog visible to both teams. No AI agent enters development without a product-approved use case entry in the shared backlog.

→ Define a handoff checklist for AI output entering the SDLC: use case approved, risk classification assigned, HITL checkpoint defined, explainability requirements confirmed.

R-02 Risk-classify every AI use case *before development.*

ADDRESSES: F-01, F-04 PRIORITY: IMMEDIATE OWNER: PRODUCT LEAD

Each AI agent or use case classified by risk level — Standard, Elevated, or High — using the Navigator Framework™ risk matrix aligned to EU AI Act categories. Risk classification determines the HITL requirement, explainability standard, and review cadence for each use case.

→ Apply Navigator Framework™ risk matrix to all agents currently in development. Retrospectively classify all deployed AI features. Document classification with reasoning for each.

→ Use classification to determine HITL requirement: Standard — named reviewer before client deployment; Elevated — named approver + review record; High — senior sign-off + documented rationale.

III. (continued)

Recommendations *(continued).*

R-03 Build the explainability framework *for EU and UK clients.*

ADDRESSES: F-03 PRIORITY: IMMEDIATE OWNER: COMPLIANCE LEAD + R&D LEAD

A structured explainability template for each AI agent serving EU or UK clients: what it does, what data it acts on, what decisions it influences, and who is accountable. Designed to meet EU AI Act transparency obligations for limited-risk systems — including disclosure, accountability documentation, and explainability for automated outputs.

→ Complete explainability documentation for all AI agents currently visible to EU/UK clients. This is the immediate regulatory remediation — existing exposure requires existing documentation.

→ Embed explainability documentation as a requirement in the R-01 integration workflow — no AI feature reaches EU/UK clients without completed explainability documentation.

R-04 Implement the HITL approval framework *for all AI use cases.*

ADDRESSES: F-01, F-04 PRIORITY: IMMEDIATE OWNER: ENGINEERING LEAD + PRODUCT LEAD

Named human-in-the-loop approval points defined for each AI use case before production deployment. A named approver reviews the AI output, confirms it meets the stated use case and is fit for client use, and signs off. The sign-off is recorded. Target: 100% HITL coverage within the four-week engagement window.

→ For each AI use case, define the named reviewer role and the approval criteria before the next deployment. No AI output reaches clients without a completed approval record.

→ Implement the approval record template: named reviewer, date, use case, confirmation statement, and any conditions noted. This record is the governance evidence — it must exist before the deployment, not be reconstructed after.

R-05 Establish the AI policy *covering all use cases.*

ADDRESSES: F-01, F-02 PRIORITY: WEEK 3 OWNER: LEADERSHIP + PRODUCT LEAD

A foundational AI use policy covering acceptable use cases, prohibited uses, data handling by AI agents, audit requirements, and the process for reviewing new AI use cases before development begins. The policy is the governance foundation — the document all other controls reference and all new AI decisions measure against.

→ Draft the policy using the framework template delivered in week two. Review with Legal and Leadership. Approve and publish to internal documentation before any new AI use case enters the development backlog.

R-06 Implement the reskilling programme *for both teams.*

ADDRESSES: F-05 PRIORITY: WEEKS 3–4 OWNER: R&D LEAD + PRODUCT LEAD

A structured capability exchange: Product team trained on AI agent capabilities, limitations, and evaluation criteria. R&D team trained on product use cases, client requirements, and the business context their agents need to serve. Outcome: both teams can brief the other effectively and jointly evaluate AI features before integration.

→ Run a two-session structured briefing: R&D to Product (AI capabilities, limitations, evaluation criteria) and Product to R&D (product use cases, client context, business requirements).

→ Establish a shared vocabulary document — agreed definitions for AI capability, risk, evaluation criteria, and governance terms — used in all joint reviews going forward.

IV.

Deliverables — *What was produced.*

D-01 AI Integration Workflow (R&D ↔ Product)

Defined workflow connecting AI development to the product SDLC. Shared backlog structure, joint review cadence, handoff checklist. Addresses F-02. Supports R-01.

DELIVERED · READY TO ACTIVATE

D-02 Risk Classification Matrix (all current AI use cases)

Navigator Framework™ risk classification applied to all AI agents currently in development and deployment. Classification documented with reasoning for each use case. Addresses F-01, F-04. Supports R-02.

DELIVERED · IMPLEMENTATION IN PROGRESS

D-03 Explainability Framework (EU/UK client-facing AI)

Explainability template completed for all AI agents currently serving EU and UK clients. Covers what the AI does, what data it acts on, what decisions it influences, and who is accountable. Addresses F-03. Supports R-03.

DELIVERED · IMMEDIATE REMEDIATION COMPLETE

D-04 HITL Approval Framework and Records Template

Named HITL approval requirements per AI use case (by risk class). Approval record template for all deployments. HITL coverage target: 100% within engagement window. Addresses F-01, F-04. Supports R-04.

DELIVERED · IMPLEMENTATION IN PROGRESS

D-05 Foundational AI Policy

Draft AI use policy covering acceptable uses, prohibited uses, data handling, audit requirements, and the review process for new use cases. For leadership review and approval. Addresses F-01. Supports R-05.

DELIVERED · AWAITING LEADERSHIP SIGN-OFF

D-06 Reskilling Programme Structure and Materials

Two-session briefing plan, shared vocabulary document, and ongoing joint review cadence. Addresses F-05. Supports R-06.

DELIVERED · SESSIONS SCHEDULED

V.

Implementation *Roadmap.*

PHASE	ACTION	TIMELINE	OWNER
Immediate	Activate R&D-to-Product integration workflow (D-01)	Before next AI deployment	Product + R&D Lead
Immediate	Apply HITL approval framework to all active AI use cases (D-04)		

		deployment	Lead
Immediate	Complete explainability documentation for all EU/UK client AI (D-03)	This week	Compliance + R&D
Week 3	Leadership review and sign-off of AI Policy (D-05)	Week 3	Leadership
Weeks 3–4	Run reskilling programme — both team briefing sessions (D-06)	Weeks 3–4	R&D + Product Lead
Ongoing	Quarterly Navigator re-calibration. Risk classification review for new use cases. Annual third-party governance assessment recommended.	Quarterly	Product Lead

Most AI governance failures aren't technical.
They're visibility failures.

This engagement has identified the gaps and delivered the governance foundation. The framework is in place. What follows is implementation — and the operational posture that means any regulator, auditor, client, or board can ask who is accountable for what the AI did, and receive a named answer with a documented record.

01 Activate all three immediate recommendations before the next AI deployment or client-facing release

02 Complete Leadership sign-off of the AI Policy (D-05) in week three

03 Run both reskilling sessions in weeks three and four — foundational for sustainable governance

04 Consider a Navigator re-calibration in six months to assess implementation and identify any new gaps as AI usage scales

Orquestra AI
rajesh@orquestraai.io
orquestraai.io