

Orchestra AI

AI DRIVES · HUMANS NAVIGATE

Two Teams. No AI Strategy. Zero Visibility.

A mid-size SaaS product and IT services company in Chennai was building AI agents into its platform. The AI R&D team and the Product team were working towards the same goal in complete isolation, with no audit trail, no governance, and no named accountability for anything the AI produced. This is what the assessment found, and what changed in four weeks.

A note on framework version. This engagement was conducted and scored under an earlier version of the Navigator methodology (three risk classes, phase-gate scoring out of 20). Current Orchestra Navigator™ engagements use the Governable Autonomy framework: a five-level maturity ladder, a five-capability assessment model, and a Certified Stage derived from evidence. The findings and outcomes below are reported exactly as delivered to the client at the time. See the Orchestra Navigator™ Framework whitepaper for the current methodology.

I

A product company accelerating into AI.

The company is a mid-size SaaS product and IT services firm based in Chennai, India, serving over 200 enterprise and SME clients across India, the UK, Australia, and the United States, across verticals including e-commerce, education, and enterprise operations. A team of under 100, moving at pace, with a product footprint larger than its headcount would suggest.

In early 2026, leadership made a deliberate decision to embed AI agents into their existing SaaS platform and develop a standalone AI agent capability for their client base. They had an R&D team already working with on-premises models. They had a product roadmap. What they did not have was any governance structure connecting the AI work to the product, the delivery process, or their customers' regulatory expectations.

The situation was compounded by a growing number of EU and UK-based clients, an audience where AI Act transparency obligations for limited-risk AI systems were already in effect. Explainability requirements, accountability documentation, and disclosure obligations were not optional for this client base. They were live.

They came to Orquestra AI with one question: "We're building AI. Are we building it right?"

Engagement structure

PHASE	ACTIVITY
Week 1 · Assessment	5-day SDLC process review, AI development workflow mapping, audit gap identification, EU compliance review.
Week 2 · Report & Workshop	Findings synthesis, risk classification per AI use case, recommendations report delivered, joint workshop with both teams.
Weeks 3–4 · Implementation	Policy drafting, HITL framework setup, explainability template development, reskilling plan initiated.

II

Two teams. No AI strategy. Zero visibility.

Before the assessment began, the expectation was that the two primary functions, the AI R&D team and the Product/SDLC team, were working in coordination. They were not. What the five-day assessment surfaced was a structural isolation that ran deeper than any process gap.

AI R&D TEAM	PRODUCT / SDLC TEAM
On-premises models · AI agent development · No product visibility · No use-case alignment	ERP product roadmap · Unaware of AI capabilities · No AI integration process · No AI evaluation criteria

Across both teams: no audit trail, no governance, no explainability for EU/UK clients.

The AI R&D team was technically proficient, experimenting with agents, fine-tuning on-premises models, building capability at pace. The Product team was managing a live multi-client SaaS roadmap. But the two teams had no shared language, no integration process, and no mechanism to evaluate whether the AI being built corresponded to the use cases the product actually served.

Neither team was failing at its own job. They were simply invisible to each other. The gap between them had no governance, no handoff, and no audit infrastructure of any kind.

"The AI team was building agents the product team didn't know existed, for use cases the product team hadn't defined, with no record that any of it had happened."

III

Five gaps. None of them technical.

Every gap identified was a governance, process, or communication failure, not a capability failure. The teams were skilled. The problems were structural.

FINDING 01 · CRITICAL · ACCOUNTABILITY

No audit trail of any kind

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 across both teams.

FINDING 02 · CRITICAL · COMPLIANCE

Complete silo between R&D and Product

The AI R&D team had no visibility into product use cases. The Product team had no visibility into AI capabilities. No integration process, shared roadmap, or handoff mechanism existed between them.

FINDING 03 · HIGH · COMPLIANCE

No explainability for EU and UK customers

EU AI Act transparency obligations for limited-risk systems applied to this client base. No such documentation existed for any AI agent in development.

FINDING 04 · HIGH · ACCOUNTABILITY

No HITL approval for AI decisions

AI agents were being developed without any defined human-in-the-loop approval points. No named person was accountable for AI outputs before they entered the product or were demonstrated to clients.

FINDING 05 · MEDIUM · COMPETENCY

Capability mismatch in both directions

The Product team lacked AI vocabulary to evaluate what R&D was building. The R&D team lacked product knowledge to prioritise the right use cases. Neither team could effectively brief the other.

IV

Six interventions. Delivered in week two.

#	INTERVENTION
01	Process bridge between R&D and Product. A defined integration workflow: shared backlog, joint use-case review cadence, handoff checkpoint before any AI output enters the SDLC.
02	Risk classification per AI use case. Each planned AI agent classified by risk level, higher-risk agents flagged for additional review before client-facing deployment.
03	Explainability framework for EU and UK clients. A structured template per AI agent: what it does, what data it acts on, what decisions it makes, who is accountable.
04	HITL approval framework. Named human-in-the-loop approval points for each AI use case before production deployment.
05	AI policy establishment. Acceptable use cases, prohibited uses, data handling by AI agents, audit requirements, and the review process for new use cases.
06	Reskilling programme for both teams. A structured capability exchange so each team could brief and evaluate the other effectively.

V

From invisible to accountable.

At the end of the four-week engagement, the company had moved from a state of zero governance to a defined, operational governance posture, with named owners, documented policies, and an integration process both teams could use from day one.

BEFORE

- Zero visibility between teams, no shared process or vocabulary
- 0 documented AI approvals for any output
- No explainability documentation for EU/UK clients
- AI use cases unclassified, no risk assessment
- No AI policy, no defined acceptable or prohibited uses
- Neither team could brief the other effectively

AFTER

- Shared R&D-to-Product integration workflow in place
- 100% HITL coverage, named approval required before deployment
- Explainability framework for all EU/UK client-facing AI
- All AI use cases risk-classified before client deployment
- Foundational AI policy established and signed off
- Both teams reskilled, shared vocabulary and joint review cadence

This engagement reflects a pattern Orquestra AI observes in organisations accelerating into AI: technical capability arrives before governance structure does. For this company, the exposure was compounded by active EU and UK client relationships. The gap was not a technical problem. It was a governance one, and it was solvable in four weeks because both teams had a common language and a clear set of named fixes.

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

Rajesh Srinivasan

Orchestra AI · rajesh@orquestraai.io
orquestraai.io

Case Study 01 · v1.1 · Distribution: Open, with permission
Details used with permission. Framework version noted on cover.
© 2026 Orchestra AI.