Orquestra AI

AI DRIVES · HUMANS NAVIGATE

Setting Governance. Getting It Right Before Scaling.

A 15-year testing services company building an AI-powered compliance checking product took it into enterprise pilots. The governance questions stopped the room. They called Orquestra AI. Two weeks. An AI Policy Framework. The product became enterprise-class, before the next pilot began.

A note on framework version. This engagement was conducted and scored under an earlier version of the Navigator methodology. Current Orquestra 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 Orquestra Navigator™ Framework whitepaper for the current methodology.

I

A testing company building a compliance product.

The company had 15 years of testing services delivery behind it, credible, technically rigorous, with a team that understood quality assurance at engineering depth. Their new product was a natural extension of that expertise: an AI-powered test automation platform with the capability to check compliance frameworks including SOX. Enterprises run the product; the AI generates findings; the customer validates and acts on them.

The product was built the way the team knew how to build things: capable engineers, coding agents, fast iteration. The compliance checks ran. The outputs looked right. When they took it into enterprise pilots, they were confident in what they had built.

They were not prepared for the questions that followed the demo.

Engagement structure

PHASE	ACTIVITY
Week 1 · Discovery	AI feature mapping, development process review, governance gap identification, risk classification per AI output type, enterprise pilot question analysis.
Week 2 · Framework Design	AI Policy Framework delivered: governance policy, customer review documentation package, AI generation controls, traceability policy benchmark, responsible AI development policy.

II

The pilot worked. The governance questions didn't.

The product demonstration went well. The AI ran the compliance checks cleanly. The output was clear. Then the enterprise security and risk team leaned in, and the room changed. Five questions. No answers available.

THE PRODUCT	ENTERPRISE PROCUREMENT ASKED
AI generates compliance findings · Customer validates before acting · SOX and other frameworks · 15 years of testing expertise	AI governance framework? · Customer review documentation? · AI generation controls? · Traceability policy?

The team had no answers. The pilot stalled, not on capability, but on governance. That is when the company called Orquestra AI.

> "The product passed every technical test. It failed every governance question. In enterprise AI procurement, those are the questions that decide the deal."

The five questions that stopped the pilot room

What is your AI governance framework? What documented policy governs how this AI operates? Our compliance team will review these findings, what documentation do they receive to validate against? What controls define the scope and criteria of what your AI checks? Your traceability controls, what policy do they enforce? How was this product built, what is your responsible AI development process?

III

Five gaps. None of them in the product.

The discovery mapped every enterprise pilot question to a specific governance gap. None of the gaps were technical failures. The product worked. The governance infrastructure around it had not been built.

FINDING 01 · CRITICAL · ACCOUNTABILITY

No documented AI governance framework

The AI operated under no documented policy. No defined rules for what it could generate, how it was supervised, or who was accountable.

FINDING 02 · CRITICAL · COMPLIANCE

Customer reviewer had nothing to validate against

Customers received findings with no documentation of the criteria, scope, or logic that produced them. Oversight without documentation is liability transfer, not oversight.

FINDING 03 · CRITICAL · CONTROLS

No documented controls on AI generation scope

No specification defined what the AI checks, how, or within what boundaries. The AI operated as an undocumented agentic engine.

FINDING 04 · HIGH · QUALITY DRIFT

Traceability controls present, no policy benchmark

Controls existed but were not effective. A trail without a policy benchmark is a log, not governance evidence.

FINDING 05 · HIGH · DEVELOPMENT

No responsible AI development policy

Built using coding agents without a spec-driven process. Same agent wrote code and tests. Test Independence Ratio: 0%.

IV

Five interventions. Delivered in week two.

#	INTERVENTION
01	**AI Governance Policy.** What the AI is permitted to generate, how it is supervised, who is accountable, how the policy is maintained. The direct answer to "what is your AI governance framework?"
02	**Customer Review Documentation Package.** Criteria applied, scope of the check, basis of the finding. Transforms customer HITL from liability transfer to meaningful governance.
03	**AI Generation Controls and Scope Definition.** What each check checks, how, on what basis, within what defined scope.
04	**Policy Benchmark for Traceability Controls.** The documented standard the existing traceability controls now enforce.
05	**Responsible AI Development Policy.** What requires human specification before coding agents begin, what requires review before shipping, traceability from requirement to code.

V

From stalled pilot to enterprise-ready.

At the end of the two-week engagement, every question the enterprise pilot room had asked had a documented, demonstrable answer. The product had not changed. The governance framework around it had, and that was precisely what enterprise procurement required to proceed.

BEFORE, COULD NOT ANSWER

- 0 governance documentation across all five dimensions
- Customer receives AI findings, no review documentation
- AI generation scope: undocumented agentic engine
- Traceability controls: present, not effective
- Development: coding agents, 0% test independence

AFTER, DOCUMENTED AND DEMONSTRABLE

- AI Governance Policy published, enterprise-ready
- Customer review package delivered with every finding
- AI generation controls documented per check type
- Policy benchmark established, controls now enforce a standard
- Responsible AI development policy implemented

This engagement reflects a pattern Orquestra AI observes in AI product companies approaching enterprise markets: technical capability arrives before governance infrastructure does. Those governance questions are not advanced requirements. They are the baseline any enterprise with its own AI risk obligations requires before deploying a third-party AI product in their environment.

AI governance cannot be an afterthought for scaling.

Rajesh Srinivasan
Orquestra AI · rajesh@orquestraai.io
orquestraai.io

Case Study 02 · v1.1 · Distribution: Open, with permission
Details used with permission. Framework version noted on cover.