PRODUCT SHEET · GOVERNANCE HUB · V1.0 MVP

Orquestra **AI** orquestraai.io April 2026

GOVERNANCE · YOU CAN SHOW

# The register, the classification, and the *evidence pack* for organisations running AI agents in production.

Register every AI agent and model as an asset. Classify it against the regulator's risk tiers. Generate the evidence pack on demand, for the auditor, the customer, the board, or the insurer. No system connections. Your data stays with you.

| AUDIENCE | STAGE | DEPLOYMENT | FAMILY |
| --- | --- | --- | --- |
| CISO · Risk · Compliance | v1.0 MVP · Live | Hosted · No integrations | Orquestra AI |

*I.*

# Your organisation is running AI in production.

*Five questions you cannot yet answer.*

In 2026, AI agents make decisions, generate code, draft customer communications, and act inside workflows across nearly every business function. Regulators have noticed. Auditors have noticed. Customers have started asking. The questions arriving on the CISO's desk are now operational, not theoretical.

QUESTION 01 · ACCOUNTABILITY

### Who owns this AI?

For every agent, model, and AI feature running in your organisation, can you name the accountable human owner, the business purpose, and the decision they personally signed off on? If the answer is "we'd have to check," you do not have an inventory.

QUESTION 02 · COMPLIANCE

### Where does the EU AI Act apply?

Which of your AI assets fall under Prohibited, High-risk, Limited-risk, or Minimal-risk categories? Can you produce the classification rationale on demand? ISO/IEC 42001, DPDP, GDPR, and sectoral regimes ask similar questions in different language.

QUESTION 03 · VISIBILITY

### Do you have a register at all?

Most organisations cannot enumerate the AI in their stack: vendor features turned on by default, agents stood up by teams, models embedded in SaaS. Without a register there is no governance, only optimism.

QUESTION 04 · QUALITY DRIFT

### Is the documentation current?

Model cards age. Risk classifications change when context changes. Vendors update capabilities silently. Without a tracked record, drift accumulates and surfaces only when something visible breaks.

QUESTION 05 · OPERATIONAL RESILIENCE

### Could you operate if these AI tools were absent, restricted, or wrong tomorrow?

Regulators increasingly expect organisations to demonstrate they have a continuity plan for AI dependencies, not just IT continuity. This is governance evidence, not architecture talk.

> *"The regulator is not asking what your AI does. The regulator is asking whether you can defend what your AI does."*

*II.*

# Three capabilities. *One governance posture.*

Governance Hub is a hosted application for organisations that need to register, classify, and evidence their AI use without exposing internal systems to a new integration. Validation is manual, by your team. No system connections. Your data stays with you.

CAPABILITY 01

### AI Asset Registry

A working inventory of every AI agent, model, vendor feature, and AI-powered workflow in scope. Each asset records named owner, business purpose, customer exposure, data classes touched, models used, and lifecycle status. Bulk-import via CSV, with no APIs to wire up.

CAPABILITY 02

### Risk Classification

Guided classification against the EU AI Act tier model (Prohibited, High, Limited, Minimal), with the rationale recorded against each asset. Surface the high-risk and customer-facing assets immediately. Re-classification on change is tracked, not silent.

CAPABILITY 03

### Evidence Pack Export

Generate a governance pack you can hand to a regulator, auditor, customer, board, or insurer. Full Export Report, Assets CSV, and the Agent-Stack Report for organisations running AI agents. The pack is named, dated, and reproducible.

## What the dashboard shows on day one

| | | | |
|---|---|---|---|
| **8** TOTAL ASSETS REGISTERED | **8** ACTIVE · 0 INACTIVE | **3** HIGH / PROHIBITED · SURFACE NOW | **3** CUSTOMER-FACING · DISCLOSURE-ELIGIBLE |
| **3** MISSING DOCS · TRIAGED FOR OWNERS | **4** DISTINCT MODELS USED | **5** STANDARDS MAPPED | **6%** COVERAGE ACROSS STANDARDS |

Coverage starts low because the dashboard tells the truth. Most organisations begin with zero structured evidence. The shape of the work, what to document, by whom, and in what order, becomes visible as soon as the register is in.

*III.*

# Aligned to the regimes already *asking for evidence.*

Governance Hub is built around the obligations regulators and standards bodies have already published. The classification model maps directly to the EU AI Act risk tiers; documentation tracking maps to ISO/IEC 42001 AI Management System clauses; and data handling is shaped to keep GDPR and DPDP exposure minimal by architecture.

## EU AI Act, risk tier model

| TIER | WHAT THE ACT EXPECTS | HOW GOVERNANCE HUB SUPPORTS IT |
|---|---|---|
| **Prohibited** | Unacceptable practices: manipulative AI, social scoring, real-time biometric ID in public spaces (with exceptions). | Flagged at classification time with a hard block on production status; rationale captured. |
| **High-risk** | AI in regulated domains: credit decisioning, hiring, critical infrastructure, medical devices, education access. The heaviest documentation burden. | Full asset record with model card, intended purpose, data classes, performance, human oversight model, and a post-market monitoring slot. |
| **Limited-risk** | Transparency obligations: chatbots, generative content, emotion recognition, deepfake disclosure. | Asset is tagged customer-facing where applicable; disclosure status tracked against each touchpoint. |
| **Minimal-risk** | The majority of AI use: productivity tools, internal automation, spam filters. | Light record, included for inventory completeness so the register is provably exhaustive, not selective. |

## Standards and regimes the platform addresses

**EU AI Act:** risk tiering, documentation, post-market monitoring records.

**ISO/IEC 42001:** AI management system clauses (scope, leadership, planning, support, operation, evaluation, improvement) mapped to evidence slots.

**NIST AI RMF:** Govern, Map, Measure, and Manage functions reflected in the asset lifecycle.

**GDPR · DPDP:** data-minimisation by design; no system connections required; data stays in your tenant.

**SOC 2 · ISO 27001:** an evidence layer that pairs cleanly with existing security attestations.

**Sectoral:** financial services (model risk management), healthcare (medical-device AI), employment (algorithmic hiring).

*IV.*

# From zero to evidence pack, *in an afternoon.*

Governance Hub is intentionally an MVP. The first version of your AI governance posture does not need a six-month programme. It needs a register, a classification, and a defensible export.

**STEP I**

### Register your assets

Upload your existing inventory via CSV, or capture assets one by one in the UI. Owner, purpose, models, data classes, and customer exposure, all captured in the same record.

**STEP II**

### Classify the risk

Walk each asset through the EU AI Act tier model. The classification is recorded with rationale and dated; re-classification on change is tracked and named.

**STEP III**

### Export your evidence pack

Generate the full Export Report, the Assets CSV, or the Agent-Stack Report for organisations running AI agents. Hand it to your regulator, customer, auditor, board, or insurer.

## What v1.0 includes

**AI Asset Registry:** CSV import, manual entry, lifecycle status.

**Risk Classification:** EU AI Act tier model, rationale, history.

**Standards module:** five standards mapped at launch.

**Governance dashboard:** coverage %, missing docs, model inventory, posture by risk level.

**Exports:** Export Report, Assets CSV, Agent-Stack Report.

**Multi-asset views:** Enterprise View, type and risk breakdowns.

**Starter Kit:** templates and reference language to accelerate adoption.

**No integrations required:** manual validation by your team; data stays with you.

## Try the Governance Hub dashboard. See where your posture actually sits.

**https://govern.orquestraai.io**

v1.0 MVP · No card · No setup call · Demo data resets in one click

Questions, a demo, or a guided walkthrough? Contact **Rajesh Srinivasan · rajesh@orquestraai.io**

Orquestra AI · orquestraai.io · April 2026 · Governance Hub · v1.0 MVP