


Every AI system, every framework. Proven with live evidence.

Regulators now ask how your AI is governed, not whether it is. Evolvable AI Risk & Compliance registers every AI system you run, classifies its risk against each framework and checks its controls continuously, from what the platform actually sees. Your posture is current every day, not once a year.

WHAT IT SEES



WHAT IT DOES

- AI systems register**
 Every AI system in one place: built on Evolvable, discovered on endpoints, or registered from other platforms. Anyone can request; your AI governance board approves.
- Risk classification**
 Intake answers suggest the risk class under each framework, such as EU AI Act high-risk. The approver decides.
- Live control checks**
 Controls verified from what the platform sees: firewall, logging, human oversight, evals and versions.
- Risk register**
 Inherent and residual risk on your own matrix. Residual risk moves as linked controls pass or fail.
- Issues & incidents**
 Track findings and incidents with regulatory clocks, containment steps and drafted reports.
- Audit-ready, always**
 Statement of Applicability, owner attestations, one-click audit pack and auditor access.

BUILT FOR

CISO

Chief Risk Officer

Data Protection Officer

AI governance lead

Model risk

Internal audit

1,400⁺

controls mapped across leading AI frameworks

15

minutes between automated control checks

1

piece of evidence, reused across every framework it covers

0

spreadsheets to chase before an audit

From intake to audit, *in one place.*

- 01

Intake

Teams request a new AI system through a short questionnaire.
- 02

Classify

Risk class suggested per framework; frameworks and controls assigned.
- 03

Assess

Impact assessments drafted by a copilot, reviewed and signed by owners.
- 04

Approve

Go-live gated until assessments cover the version that will run.
- 05

Monitor

Live checks, issues and incidents. A model or tool change reopens review.
- 06

Audit

Evidence, attestations and a one-click audit pack, always current.

EVIDENCE YOU CAN TRUST

Every control shows *how sure we are.*

Declared evidence alone can't prove a technical control works. The module shows the level each control has reached and what it takes to move up.

Observed

Runtime traffic through Evolvable or the AI gateway. Controls verified every 15 minutes.

Tested

Results imported from red-team and testing tools, such as Garak, PyRIT, Promptfoo or a partner.

Declared

Questionnaires, documents and owner attestations.

VERSION-AWARE GOVERNANCE

Change the agent, *reopen the review.*

When a new version changes the model or adds a tool, the approval and impact assessment are reopened automatically. High-risk systems can't ship until the paperwork covers the version that will run.

Fix it in one click.

A failing control links to the setting that fixes it, with a preview of what changes. Controls re-check as soon as you apply it.

ONE AI SYSTEM, EVERY FRAMEWORK Example record. Illustrative.

Credit Pre-screen Agent Retail Lending · Owner: Head of Retail Credit · Built on Evolvable · v14 HIGH-RISK		
FRAMEWORK	HOW IT CLASSIFIES THIS SYSTEM	CONTROLS
EU AI Act	High-risk · credit decisions about individuals	31 pass 4 fail
ISO/IEC 42001	Impact assessment required	22 pass 2 fail
DIFC Regulation 10	Autonomous system making decisions on individuals	14 pass 1 fail
NIST AI RMF	Likelihood × magnitude mapped	9 pass 1 partial

Top fix: the decline action runs without a person deciding. Add an approval gate to close 3 failing controls at once.

FRAMEWORK LIBRARY Turn on the frameworks that apply. Scoping questions keep out obligations that don't.

- EU AI Act
- ISO/IEC 42001
- NIST AI RMF
- NIST AI 600-1
- OWASP Top 10 for LLMs
- OWASP Agentic
- DIFC DP Regulation 10
- UAE FI Enabling Technologies
- SDAIA (KSA)
- Your internal AI policies

BUILT INTO THE PLATFORM

- AI Firewall

Policies show the controls they satisfy. Turning one off warns you first if it would fail a control.
- Every agent

A compliance tab on each agent shows its failing controls and how to fix them.
- Human oversight

Approvals and sign-offs use the same oversight groups that approve agent actions.
- Evolvable Edge

Shadow AI seen on endpoints is registered, assessed and covered by policy.

See your AI compliance posture on your own systems.

Pick one framework and the AI systems you already run. We register them, run the checks and show you where you stand, with the evidence behind every score.