DATASHEET

Archer Evolv™ AI Compliance

Compliant AI and agentic workflows, from regulation to runtime, with the lineage to prove it.

Archer Evolv™ AI Compliance makes your AI and agentic workflows compliant with lineage, runtime governance and issue management: regulation and policy become working controls, enforced pre-inference and evidenced on a trusted system of record.

The Challenge

AI adoption compounds. In 2026, AI became the second-highest ranked global business risk, its biggest jump on record,[1] and governance and compliance still advance one control at a time. AI agents now perform tasks previously managed by employees and third parties on the enterprise's behalf, and can unknowingly expose sensitive data or take actions that conflict with regulation and company policy. Traditional GRC programs were designed around employees and third parties, not autonomous AI agents, and that gap lands on four different teams:

- **Compliance.** Obligations arrive faster than manual enforcement can translate them: the gap Compliance & Governance teams live in every day.
- **Output & data.** Unsafe answers, data where it should not go: exactly what Risk teams are on the hook to validate.
- **Shadow.** AI nobody sanctioned, invisible and unowned: what Security teams can't see until it's a problem.
- **Bypass.** Sanctioned AI skipping the governed path: the gap AI & Platform teams are asked to close.

Closing them takes four things: lineage to the regulation, prevention before the model, expert-in-the-loop approval, and one examiner-ready record. Measure any solution against this list.

EU AI Act: live

Article 50 transparency duties and enforcement powers in force since 2 August 2026; penalty ceiling €35M or 7% of global turnover.

AI ACT · ART. 99 · PENDING DIGITAL OMNIBUS REVISIONS

78%

Of executives lack strong confidence they could pass an independent AI governance audit within 90 days.

GRANT THORNTON · 2026

1. Allianz Risk Barometer, 2026.

Overview

Archer Evolv™ AI Compliance makes your AI and agentic workflows compliant with lineage, runtime governance and issue management. Regulation and policy become working controls (regulation in, guardrails out, evidence back), enforced before the model is invoked and evidenced on a trusted system of record.

Purpose-built AI operators do the work, each a task-specific digital worker that performs one job under supervision: identity-bound, scope-constrained, expert-supervised, audit-emitting. Enforcement is native: controls deploy as AWS Bedrock guardrails inside your own AWS account, with no proxy in the inference path and no added latency. Models running outside Bedrock are covered by the same approved control. Behind them, the GRC record 1,300 organizations run on, including half the Fortune 500.

CAPABILITIES	Regulations→Controls · Controls→Guardrails · Guardrail conformance · Violation monitoring · Issue management
RUNS NATIVELY ON	AWS Bedrock; models running outside Bedrock are covered by the same approved control through ApplyGuardrail
EXPOSED TO AGENTS	callable by any agentic workflow, copilot or pipeline over MCP, REST and webhooks

Every obligation becomes a tracked control. Runtime-enforceable ones deploy as guardrails that stop violations before the model is invoked; the rest are worked as controls on the same record. Nothing enforces without your approval, and every step is reproducible for an examiner with full lineage.

How It Works

01	**LISTEN** **Sources → Controls**	9 tier-1 AI regulations, hundreds of adjacent data & privacy sources, and your own policies become tracked controls.
02	**DECIDE** **Controls → Guardrails**	Enforceable controls become draft AWS Bedrock guardrails to run in observe mode on live traffic, deployed on your approval.
03	**ACT** **Violation monitoring**	Traffic is evaluated against the guardrail pre-inference in your AWS account; violations are blocked and logged with control, policy and confidence.
04	**ASSURE** **Guardrail conformance**	Deployed guardrails tested on cycle against the control you approved: drift and tampering flagged first.
05	**LEARN** **Issue management**	Findings route into Archer issue management workflows, owned and driven to closure on a trusted system of record.

Tunable enforcement: observe logs only, advise notifies an owner, enforce discards pre-inference. Every step immutably versioned. Reproducible on demand. Audit-ready.

How It's Delivered

Getting started costs minutes, not a program. Your administrator creates one scoped, least-privilege AWS role; the AI operators assess your live estate; and your executive report (what's running, what's governed, what's exposed) is reviewed together on day ten. Nothing is deployed until you approve.

DAY 0
Your admin, one scoped role

DAYS 1-5
We onboard, read-only assessment

DAY 10
Executive report, reviewed together

DAYS 11-20
First guardrails, on your approval

ONGOING
Steady state, at your cadence

Delivered as a Turnkey Solution

Your Archer team runs the program with you through the AI operators. Onboard once, then assess, review and tune at the cadence you choose, quarterly through weekly. You approve; we do the work.

Onboard once

- Connect your account
- Build your profile
- Load regulation & policy
- Stand up guardrails

Assess

- Run conformance checks
- Report violations & drift
- Issue the readout
- Prioritize actions

Review

- Walk the readout
- Agree what matters now
- Set next priorities
- Track progress

Tune

- Refine your profile
- Update controls & scope
- Promote tuned guardrails
- Deploy on your approval

Built for Every Team That Owns the Outcome

The CISO owns the platform decision and the budget. The evaluation itself usually runs inside the security organization, with compliance, risk, platform and audit weighing in alongside.

Security Teams

Keep data secure when AI, not a person, is the one acting.

Every action attributable in your own audit logs, through a role that reads and suggests but never deploys.

Compliance & Governance

Show the control, not just the policy.

Every violation and control tied to an obligation, with the full lineage reproducible whenever an examiner asks.

Risk Teams

Report one risk posture, not five dashboards that disagree.

Every invocation evaluated against the controls that bound it, and blocked pre-inference when it falls outside them.

AI & Platform Teams

Ship compliant AI faster, with no console to fight.

Connect once, no proxy, nothing to run. One control set across your model catalog, every model and agent through the same gate.

Executive Leadership

Move faster, and more responsibly.

A straight answer, any time, to what's running, what's governed and what's exposed, so each new use case ships without a bespoke review.

Internal Audit

Sample AI decisions the way you sample a human reviewer's.

Every AI decision carries the same evidence and lineage as a human one, so nothing needs a special request to test.

Features and Benefits

Feature	Benefit
Regulatory intelligence: Tier-1 AI regulations, hundreds of adjacent data and privacy sources and your policies, monitored continuously for change.	Nothing missed when the rules change, and far less regulation read and configured by hand.
Runtime enforcement: enforceable controls deployed as AWS Bedrock guardrails in your own AWS account, one set across the model catalog.	Violations blocked pre-inference, natively in your AWS account, consistently across every model.
Model & platform coverage: one approved control set spans the leading foundation models in Bedrock, and reaches models running outside Bedrock through the same control.	One policy enforced across the estate, with no per-model drift and no second toolchain.
Governed change: tunable enforcement (observe, advise, enforce) with approval gates and immutable versioning.	Nothing enforces without your approval, and any version rolls back instantly.
Evidence and lineage: the full chain from authoritative source to obligation to control to guardrail to violation event, immutably versioned and reproducible.	Examiners get proof, not attestation, whenever they ask for it.
Issue management: violations, drift and findings open as issues, routed to named owners on a trusted system of record, with status and closure tracked.	Findings get worked and closed, not just logged, on the same lifecycle as the rest of GRC.
Agentic interoperability: AI compliance exposed to any enterprise agentic workflow, copilot or pipeline over MCP, REST and webhooks.	Consistent enforcement across teams, tools and agents, with no drift between them.
Boundary and accountability: guardrails run inside your own account; Archer receives the violation event only.	Prompts stay inside your boundary, and AI compliance sits on the same trusted system of record.

Preventing a violation and proving why it was prevented are two different problems: every other category solves one.

For your CCO, the exam question has changed: from “show us your policy” to “show us the control.”

Start With Your Read-Only Assessment

Contact your Archer account team. Product detail, implementation guides and platform setup walkthroughs at ArcherIRM.com, or scan to go straight there.

QR CODE

About Archer

For more than 25 years, Archer has pioneered holistic governance, risk and compliance (GRC) solutions that give enterprise organizations a clear understanding of risk.

Visit www.ArcherIRM.com **@ArcherIRM** **Archer Integrated Risk Management**

©2026 Archer Technologies LLC. All rights reserved. Archer and the Archer logo are registered trademarks or trademarks of Archer Technologies LLC in the United States and other countries. All other trademarks are the property of their respective owners. Archer believes the information in this document is accurate. The information is subject to change without notice. 8/26 Data Sheet.

ARCHER®