ARCHER®

SOLUTION BRIEF

Archer Evolv™ AI Compliance

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

Archer turns the regulations that apply to you, together with your own AI and data-privacy policies, into live AWS Bedrock guardrails, enforced before the model is invoked and evidenced on the record your auditors already trust.

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 governance gap shows up as four risks. Each one raises a question you should already be able to answer.

- **Compliance.** Regulation lands faster than any manual process can translate it, the gap your compliance team lives in every day. *Which obligations actually apply?*
- **Output & data.** Records, contracts and source code leave with every prompt, exactly what your risk team is on the hook to validate. *What reaches the model, and what comes back out?*
- **Shadow.** AI runs where no control was ever attached, invisible to security until it is a problem. *What is running that nobody approved?*
- **Bypass.** Sanctioned AI slips off the governed path, the gap your AI and platform team is asked to close. *What has drifted since you approved it?*

Closing the gap requires a solution that does four things: prevent before the model, create lineage to the regulation, keep an expert in the loop, and hold one examiner-ready record.

52% / 90%

Workers using unapproved AI vs. executives sure they'd see it.

OKTA · 2026

1 in 3

Office professionals have put customer data into public AI.

PAGERDUTY · 2026

€35M / 7%

EU AI Act maximum penalty, whichever is higher.

AI ACT · ART. 99 · PENDING REVISIONS

€2.4M

Average GDPR fine; cumulative penalties now past €7.1B.

DLA PIPER · JAN 2026

1. Allianz Risk Barometer, 2026.

Our Solution

Regulation in. Guardrails out. Evidence back.

Archer Evolv™ AI Compliance makes your AI and agentic workflows compliant from regulation to runtime, with the lineage to prove it. Regulation and policy become working controls, enforced pre-inference and evidenced on the Archer system of record. AI operators do the work: purpose-built digital workers, each performing one job under supervision. 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.

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 the Archer system of record.

The runtime-enforceable subset becomes guardrails; the rest are worked as controls on the same record.

Everyone Does One. We Do Both.

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

WATCHES, FINDS IT LATER

Sees it. Can't stop it.

Governance platforms · observability · DLP

"We detect violations."
After the output already shipped.

"We map you to frameworks."
But nothing enforces the mapping.

BLOCKS, CAN'T EXPLAIN IT

Stops it. Can't prove it.

Guardrail firewalls · prompt filters

"We block bad prompts."
But not which rule required the block.

"We log everything."
Logs aren't lineage.

PREVENTS AND PROVES IT

Both halves, one chain.

Archer Evolv™ AI Compliance

- **Prevention.** Violations blocked pre-inference
- **Lineage.** Every block traced to the obligation behind it
- **Approval.** Nothing enforces without your approval
- **One record.** Source to event, exported on demand

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

The Architecture

The layer AI & Platform teams connect, Security teams trust, and Compliance & Governance teams draw evidence from. It runs natively on AWS Bedrock: guardrails deploy inside your own AWS account, and models running outside Bedrock are covered by the same approved control. The same compliance layer is exposed to agents: any agentic workflow, copilot or pipeline can call it over MCP, REST and webhooks.

Authoritative Sources

- AI regulations
- Frameworks
- Data & privacy regulations
- Your policies

Archer Platform

- Business profile
- AI operators
- Archer system of record: controls, evidence and issues

Your AWS Account

- Users & applications
- AWS Bedrock guardrails
- Foundation models
- Invocation logs

Archer sets the guardrails; AWS Bedrock runs them, with no proxy in the inference path. Identity-bound, scope-constrained AI operators run the loop; you approve, and every action lands on the record.

BUILT ON ARCHER'S PATENTED TECHNOLOGY

18	492	130+	8,000+	1,300	half
patents	purpose-built models	GRC experts	regulatory & standard-setting sources	organizations	of the Fortune 500

Tunable Enforcement

Enforcement is a dial you hold. Nothing blocks until you promote it.

Observe.

Observe mode on live traffic. Logs what it would block; blocks nothing.

Advise.

Findings raised to a named owner with evidence attached. Nothing enforced.

Enforce.

Violations blocked pre-inference.

Every promotion up the dial requires your approval. Every version is immutably recorded and rolls back instantly.

Your Prompts Never Leave Your Account

Archer connects through one scoped, least-privilege AWS IAM role that reads your guardrail configuration and the events it produces. It proposes and evidences; it does not read your traffic.

NEVER REACHES ARCHER

- Prompt content and model responses
- Documents, embeddings and vector stores
- Customer PII
- Model weights and training data

WHAT ARCHER RECEIVES

- That a violation occurred
- Which control fired, for whom and when
- Confidence score and policy type
- Guardrail configuration and version history

Read-and-suggest only: nothing deploys without your approval. If Archer is unreachable, your guardrails keep enforcing exactly as last deployed.

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.
Remediation workflow: violations, drift and findings open as issues, routed to named owners on the Archer 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.

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.

How It's Delivered

From First Access to First Report

Getting started costs minutes, not a program. Your administrator creates the scoped AWS IAM 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	DAYS 1-5	DAY 10	DAYS 11-20	ONGOING
Your admin, one scoped role	We onboard, read-only assessment	Executive report, reviewed together	First guardrails, on your approval	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 approval

Start With Your Read-Only Assessment

Contact Archer, or your account team if you are already a client. 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**

ARCHER®

©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 Solution Brief.