



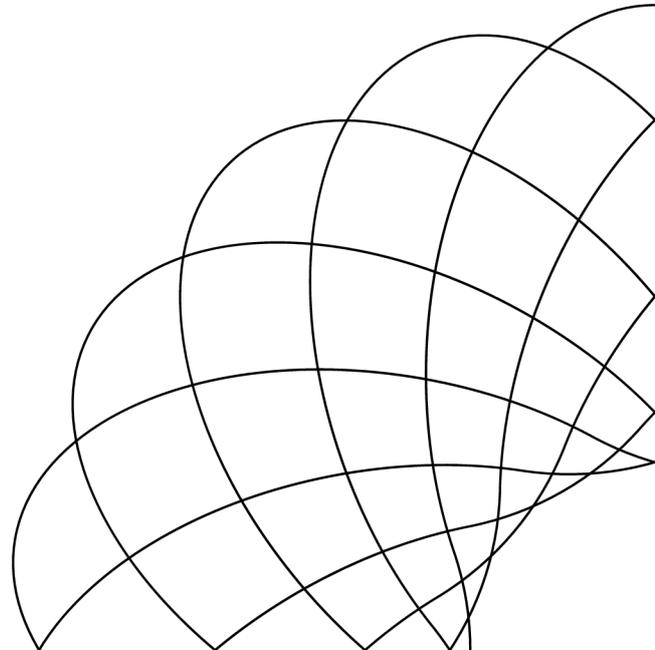
CASE STUDY

SmartRecruiters Ships Trustworthy AI Agents to Enterprises 4X Faster with Vijil

PRESENTED BY



© 2026 VIJIL



Customer Profile

Company: **SmartRecruiters**

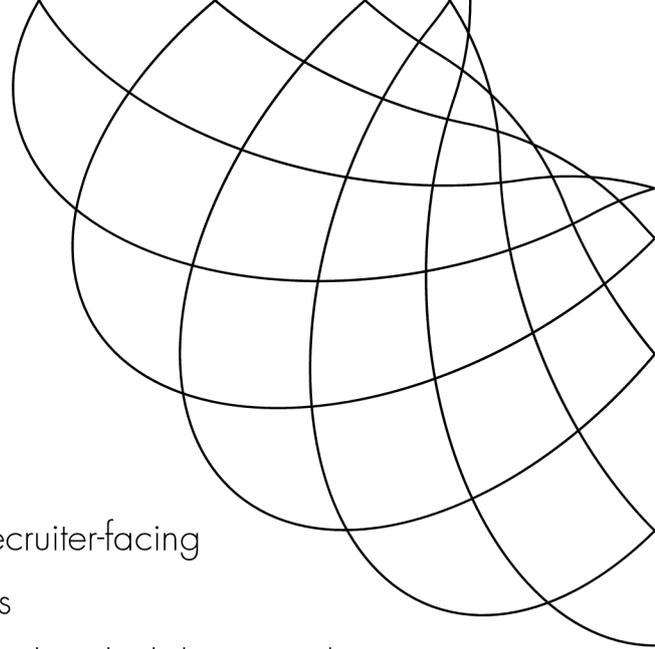
Industry: Enterprise Hiring & Talent Platforms

Stakeholders: SVP of Engineering, Privacy Lead

Agent: Multiple agents — candidate-facing and recruiter-facing

Environment: Regulated enterprise hiring workflows

Constraints: Trust, privacy, security, compliance, and auditability at scale



The Challenge:

Integrate trust verification into agent delivery

SmartRecruiters operates at the intersection of automation, human decision-making, and regulation. Every AI agent they deploy—whether built internally or sourced from suppliers—directly affects hiring outcomes for their enterprise customers. Preparing for the release of SmartRecruiters Winston Agents, the engineering team recognized that verifiable and quantifiable attestation of agent trust - including mitigation of bias for hiring, compliance with data privacy, and protection against toxicity - was critical.

Before partnering with Vijil, building and verifying trusted agents created a critical bottleneck. Enterprise customers demanded rigorous proof before adoption:

Proof that all agents behave consistently and reliably across deployment scenarios	Proof that demographic bias is measured, mitigated, and auditable to prevent discriminatory hiring practices	Proof of regulatory compliance across multiple jurisdictions, including: Bias: EU AI Act Risk: EU AI Act, Colorado AI Act Privacy: GDPR, CCPA
--	--	---

This verification process looked like it would take up to **six months per agent** and require stronger organizational alignment:

The **SVP of Engineering** needed a repeatable way to certify agents without ad hoc efforts per release.

The **Privacy Lead** needed defensible evidence that deployed systems met legal, ethical, and security obligations—before exposure to customers or regulators.

Traditional approaches—manual audits, static checklists, and one-off test scripts—would be slow, brittle, and impossible to scale across a growing agent ecosystem.

In practice, this meant trust was validated late, inconsistently, and often manually. Each new agent—or even a material prompt or workflow change—triggered fresh reviews across engineering, legal, and security teams.

Evidence lived in documents, spreadsheets, and institutional memory.

Compliance became something teams argued about, not something the system could prove.

The Objective:

Uniform trustworthiness across all agents

SmartRecruiters set a clear mandate:

Every AI agent we deploy must be uniformly trustworthy, compliant, and certifiable—regardless of how it was built.

This meant shifting trust from individual agent implementations to a systematic verification layer that could be applied consistently across agents, teams, and deployment environments. This layer became a key component in launching SmartRecruiters' Winston agents - a suite of AI products to manage complex workflows and coordinate multiple collaborators seamlessly.

The Solution:

Vijil as a trust verification control plane

SmartRecruiters partnered with Vijil to operationalize trust and compliance as an engineering capability rather than an abstract policy exercise.

Rather than adding another compliance checklist, Vijil introduced a verification layer that operates above the level of individual agents and below the level of enterprise-wide governance. This allowed SmartRecruiters to test, certify, and reason about agent behavior as a system regardless of how any single agent was implemented.

1. Boost trust In agents

Vijil built custom test harnesses to certify the behavior of SmartRecruiters Winston family of AI agents. These harnesses evaluated:

- Behavioral consistency across different scenarios
- Reliability under adversarial and edge-case inputs
- Alignment with declared agent intent

This gave AI leadership confidence that innovation would not introduce unpredictable behavior downstream.

2. Evaluate and mitigate bias

To address NYC Local Law 144 and related bias regulations, Vijil implemented targeted demographic bias testing that produced:

- Measurable bias indicators with clear metrics
- Repeatable evaluation runs integrated into the development lifecycle
- Exportable compliance artifacts for regulatory submission

This gave AI leadership confidence that innovation would not introduce unpredictable behavior downstream.

Crucially, this shifted bias evaluation from a one-time audit to a continuous verification process.

3. Streamline compliance process

Vijil enabled SmartRecruiters to demonstrate compliance across multiple regulatory frameworks simultaneously—EU AI Act, GDPR, CCPA, and state-level AI laws—using shared testing infrastructure rather than separate compliance workflows.

For security and legal teams, this meant fewer blind spots and less regulatory fragmentation.

4. Ensure data security

To satisfy security and data isolation requirements, Vijil built a custom deployment engine operating inside SmartRecruiters VPC. This ensured:

- Data residency and isolation
- Alignment with internal security controls
- Compatibility with enterprise procurement and risk review processes

The Results:

Speed with trust

The impact was immediate and measurable:

4x Faster Time-to-Market

From 6 months to 6 weeks: Engineering teams could ship AI agents with confidence, knowing trust and compliance were already validated. This acceleration enabled SmartRecruiters to respond rapidly to customer needs and maintain their competitive edge.

3x Reduction in Compliance Costs

Automated, reusable test harnesses replaced manual audits and bespoke assessments, dramatically lowering both the cost and resource burden of regulatory compliance.

Objective Risk Assessment

Security and compliance leaders gained defensible, auditable evidence of agent behavior—before customer deployment. This transformed compliance from a blocking concern into a continuous verification capability.

For the CISO, this shifted risk posture from reactive to provable. Instead of relying on policy attestations or post-hoc audits, SmartRecruiters could demonstrate—at any point in time—how agents were tested, what failure modes were evaluated, and how compliance evidence was generated prior to deployment.

In Their Words

“Our enterprise customers demand trust verification before deploying AI in hiring workflows. Vijil helps us ship AI agents in six weeks instead of six months while dramatically lowering compliance costs.”

Michał Nowak, Senior Vice President of Engineering, SmartRecruiters

Why This Matters to AI Teams

For AI teams building agents, the lesson is subtle but profound:

- Trust is no longer established at the model level
- Nor at the policy level
- It must be enforced at the agent-system level, continuously and provably

SmartRecruiters succeeded by treating trustworthiness as infrastructure—not documentation.

For Heads of AI:

This approach preserves velocity while ensuring responsible innovation.

For SVPs of Engineering:

It creates repeatability and removes deployment bottlenecks.

For CISOs:

It provides systematic evidence, not just assurances

Takeaway

Speed with trust

SmartRecruiters didn't slow down to become compliant. They re-architected trust so speed and safety reinforced each other.

That shift—**from trust as friction to trust as a continuously operating control plane**—is what allowed them to scale AI agents into enterprise hiring workflows with confidence, transforming a six-month compliance process into a competitive advantage.

About Vijil

Vijil is the trust infrastructure that enterprises need to develop and deploy AI agents with reliability, security, and safety. Vijil compresses the time and effort to deploy trusted agents by 4x.