

AI MANAGEMENT SYSTEM · OPERATIONALIZED

ISO 42001 Compliance: The path to operationalizing agentic AI

Built on the lifecycle the standard already describes.

Vijil's capabilities map tightly to the technical requirements of the ISO 42001 lifecycle to close the agentic gap with evidence — know what you have, evaluate it before deployment, assess and understand risk as part of a governance process, protect it in production, and improve continuously based on what actually happens in production.

Vijil's Trust Score generated from bespoke and adaptive testing of an agent for security, reliability, and safety serves as both as a benchmark throughout the lifecycle, and the evidence that auditors need to evaluate for agent trustworthiness assessment.

THE LIFECYCLE, CONTINUOUSLY

- 01 **Know** A.3 / A.4
Discover & register every agent
- 02 **Evaluate** A.5 / A.6.3
Bespoke testing for trustworthiness, including adaptive, multi-turn adversarial
- 03 **Protect** A.6.5
Policy-driven guardrails that run as part of the agent, highly performant
- 04 **Improve** A.6.5
Adapt agents based on real-world behavior with remediation proposals for developers

FINDINGS FEED BACK INTO EVALUATION

WHY ORGANIZATIONS HAVE EMBRACED ISO/IEC 42001

REGULATORY ALIGNMENT

It is the certifiable backbone regulators point to. Conformance maps directly onto EU AI Act obligations, giving one structure to satisfy many jurisdictions.

PROCUREMENT LEVERAGE

Enterprise buyers and public bodies increasingly require it. Certification shortens security reviews and removes AI governance as a deal blocker.

BOARD-LEVEL ASSURANCE

Clauses 5 and 9 force named accountability and management review, so AI risk reaches the board as measured evidence rather than anecdote.

SHIPPING VELOCITY

A defined risk and control process replaces case-by-case debate, so approving each new agent stops being a bespoke negotiation.

THIRD-PARTY COVERAGE

Annex A.10 extends governance to models and agents consumed through vendors and APIs — where most real exposure now sits.

ONE SYSTEM, MANY FRAMEWORKS

It reuses existing ISO 27001 and SOC 2 machinery, so AI governance extends a working management system instead of starting another one.

WHAT OPERATIONALIZING LOOKS LIKE WITH VIJIL

01 LIFECYCLE

Mapped to the clauses, not around them

Capabilities line up to Clauses 4–10 and Annex A.2–A.10, so evaluation scope, controls, and review records extend an existing AIMS instead of forming a parallel program.

02 EVIDENCE

Every run leaves an artifact

Dated test reports, coverage matrices, guardrail logs, and a reproducible Trust Score — generated whenever a system changes, not reconstructed before an audit.

03 CONTINUOUS

Compliance as infrastructure

Continuous testing, monitoring, and improvement run as part of how agents ship. Certification becomes a byproduct of the pipeline rather than a scramble every three years.

WHY VIJIL FOR ISO 42001

One trust layer for the whole agent lifecycle.

Discovery, registration, evaluation, protection, and monitoring each map to a specific ISO 42001 annex control — and each produces a specific artifact an auditor can check.

DISCOVER ANNEX A.4	REGISTER ANNEX A.3	EVALUATE ANNEX A.6.3	PROTECT ANNEX A.6.5	MONITOR ANNEX A.6.5	ADAPT ANNEX A.6.5
-----------------------	-----------------------	-------------------------	------------------------	------------------------	----------------------

STAGE	CONTROL	WHAT VIJIL DOES	AUDIT EVIDENCE PRODUCED
Discover	A.4	Continuously scans repos and cloud/Kubernetes environments to find every agent — including shadow AI never formally onboarded.	A centralized, always-current system inventory that becomes the authoritative “in-scope” list everything else depends on.
Register	A.3	Assigns each governed agent a persistent identity, owner, and documented policy scope.	A system-generated ownership record answering “who owns this system” — not an org-chart assumption.
Evaluate	A.6.3	Diamond runs bespoke tests across reliability, security, and safety using the 3Ps method (persona, policy, purpose); adaptive offensive testing runs autonomous multi-turn attacks mapped to the OWASP ASI Top 10.	Repeatable, dated test reports and a coverage report showing exactly which risk categories were tested, and how deeply.
Trust Score	A.5	Converts evaluation results into one reproducible score tied to specific risk-register dimensions.	A dated, versioned record proving a risk-register entry was empirically tested — not just discussed in a workshop.
Protect	A.6.5	Dome enforces the same policies at runtime, filtering unsafe inputs and outputs and flagging anomalies.	Dashboards, alert history, and policy-enforcement logs proving governance runs continuously in production.
Adapt	A.6.5	Darwin analyzes production telemetry — edge cases, failures, drift — and feeds fixes back into the system.	Root-cause analyses and remediation records showing the AIMS is reassessed as systems change, not filed once and left static.

ANNEX A COVERAGE AT A GLANCE

A.2	AI policy enforced via Dome guardrails	SUPPORTING	A.3	Internal organization — Console ownership & visibility	SUPPORTING
A.4	Resources for AI systems — Platform documents tooling	SUPPORTING	A.5	Assessing impacts — Diamond scores reliability, security, safety	DIRECT
A.6	Full lifecycle — verification, operation, monitoring, logging	DIRECT	A.7	Data for AI systems — leakage/PII testing, in-VPC deployment	SUPPORTING
A.8	Information for interested parties — Trust Scores & reports	DIRECT	A.9 / A.10	Use of AI & third-party relationships — runtime + vendor coverage	DIRECT

CUSTOMER PROOF · SMARTRECRUITERS

From reactive risk posture to provable evidence, for a live hiring agent.

Winston, SmartRecruiters’ hiring agent family, was already tested against the EU AI Act, GDPR, CCPA, and NYC Local Law 144. Vijil’s behavioral, bias, compliance-mapping, and data-security evaluations became the foundation for operationalizing ISO 42001 on top of that same evidence base.

6wk TIME-TO-MARKET, DOWN FROM 6 MONTHS
3x LOWER COMPLIANCE COST VIA REUSABLE HARNESSSES

See your ISO 42001 evidence generated automatically.

Book a 30-minute walkthrough of Vijil across the agent lifecycle.

vijil.ai · sales@vijil.ai