Conformant

Compliance-as-code SDK specifically for AI-product obligations — EU AI Act risk classification, model cards, training-data provenance logs, plus the usual SOC2/HIPAA hooks. Vanta exists but isn't built for the AI-specific layer.

Category	Set 3 · Post-AI Plays
Customer	AI-product startups + mid-size enterprise AI teams shipping LLM-powered features subject to AI-specific compliance obligations
Monetisation	$1.5k–10k/month SaaS by company stage · $25k–80k/year enterprise · audit-prep engagements priced separately
Build effort	Med
Plan version	v1.0 — 2026-05

Executive Summary

Conformant is a compliance-as-code SDK and SaaS platform for AI-product companies facing the new wave of AI-specific compliance obligations. The EU AI Act entered force in 2024 with provisions phasing in through 2026-2027. New York City Local Law 144, Colorado SB 205, the proposed US federal AI safety frameworks, ISO/IEC 42001 — each introduces specific requirements for AI products that the generalist compliance tools (Vanta, Drata, Secureframe, Tugboat Logic) do not address adequately. AI companies are stuck between expensive compliance consultants ($120k-400k engagements) and unstructured DIY effort that produces audit-fail risk.

Conformant addresses the AI-specific layer that sits alongside the generalist SOC2/HIPAA/ISO 27001 hooks. Core capabilities: EU AI Act risk-classification workflow producing the required risk-management documentation, model-card generation tied to actual model deployment, training-data provenance logging with version control, intended-use documentation that surfaces in-product, post-market monitoring integration, bias-and-fairness testing workflow, incident-reporting automation for serious AI incidents. Pricing is $1.5k-10k/month SaaS tiered by company stage (early-stage AI startup, mid-stage Series B+, enterprise).

Year-1 target: 180 paying subscribers + 25 enterprise customers, generating $2.8 million revenue (~■23 crore) against ■9.5 crore in costs. The wedge against generalist compliance platforms is AI-specific depth (Vanta's 'AI policy template' is not the same as an EU AI Act risk-classification workflow with auditable artefacts). The wedge against compliance consultants is the cost structure: $30k/year ARR per customer vs. $120k+ per consultant engagement.

The Problem

AI companies in 2026 face a compliance landscape that the generalist tools were not built for. The EU AI Act (phased through 2026-2027) requires risk classification of every AI system, technical documentation, training-data summaries, post-market monitoring, fundamental rights impact assessments for certain use cases. NYC Local Law 144 requires bias audits for automated employment decision tools. Colorado SB 205 requires impact assessments and consumer disclosures for high-risk AI systems. ISO/IEC 42001 provides an AI management system standard that some enterprise customers are beginning to require from vendors.

The traditional compliance platforms (Vanta, Drata, Secureframe, Tugboat Logic) automate SOC2, ISO 27001, HIPAA, GDPR — the established compliance regimes. They have begun adding AI-related templates and questionnaires but their depth is shallow: a generic 'AI policy' template, a vague 'AI risk assessment' workflow, no integration with actual model deployments or training-data pipelines. An AI company using Vanta for SOC2 still needs to do its EU AI Act work essentially from scratch.

The alternative is compliance consultants. Specialty AI compliance consultancies (BSI's AI Trust Practice, BABL.ai, DLA Piper's AI practice) charge $120k-400k per engagement for AI Act readiness work. For an early-stage AI startup at $4-8M ARR, this is prohibitive. They are stuck between expensive consulting and unstructured DIY effort that produces audit-fail risk.

The market gap: a focused, AI-specific compliance platform that addresses the AI-Act + adjacent obligations at SaaS pricing ($18k-120k annually) with audit-ready artefacts that pass scrutiny. This is the gap Conformant fills.

The Solution

Conformant provides AI-specific compliance-as-code organised around the actual obligations AI companies face. Core capabilities: (1) EU AI Act risk classification — guided workflow that walks the company's AI system through the Act's risk categorisation (prohibited, high-risk, limited risk, minimal risk) with documented reasoning and supporting evidence. (2) Technical documentation generation — automated production of the required technical documentation under Annex IV, with placeholders the company fills via structured prompts and automated population from model-card and training-data logs. (3) Training-data provenance logging — SDK and tracker that logs training-dataset versions, sources, licensing, transformations; produces the data-governance summaries Article 10 requires. (4) Model card generation — tied to actual model deployments; auto-updates when models are retrained or replaced.

Additional capabilities. (5) Intended-use documentation — surfaces in-product where required (transparency obligations for limited-risk systems), tracked for currency. (6) Post-market monitoring — integration with production observability tools (Helicone, LangSmith, Datadog) to capture the model-performance and incident data the AI Act requires for high-risk systems. (7) Bias-and-fairness testing — workflow for running protected-class-attribute bias testing with documented results, particularly relevant for NYC Local Law 144 and Colorado SB 205. (8) Incident reporting — automated workflow for serious-incident reporting under the AI Act and adjacent regulations.

Three structural differences from generalist platforms define the wedge. First, AI-specific depth: every capability is built around the actual obligations of the AI-specific regulations rather than generic security/privacy frameworks. Second, technical integration: SDKs and APIs that integrate with the company's actual ML / LLM stack (the training-data logger plugs into HuggingFace datasets, MLflow, Weights & Biases; the model-card generator pulls from model registries; the monitoring integration plugs into Helicone, LangSmith, OpenTelemetry). Third, audit-ready artefact production: every workflow produces structured documents and audit trails designed to be presented to regulators or notified bodies — not just internal templates.

Market Opportunity

Addressable AI-company market in 2026: approximately 8,500 companies globally shipping LLM-powered products or features at sufficient scale to fall under AI Act high-risk categorisation or adjacent regulations. Of these, an estimated 3,500-4,500 are willing-to-pay at Conformant's SaaS pricing.

At a blended ARPU of $28,000/year per subscriber (across the tiers), the SAM is approximately $100-130 million growing at 60%+ annually as AI Act enforcement phases in and additional regulations (US federal AI framework, additional state laws, sector-specific AI regulations) come into force.

Realistic capture: 2-4% of SAM by year 3 = $3-5M ARR. Comparable trajectory to Vanta in its first 3 years (Vanta scaled to $35M ARR in year 3, but in a larger and more mature market; Conformant's market is forming as Vanta's was).

Adjacent expansion opportunities: regulatory consulting services (Conformant's structured artefacts make it possible to offer fixed-price audit-prep engagements at $25-80k that compete with the $120k+ traditional consultants), insurance product partnerships (AI errors-and-omissions insurance pricing benefits from Conformant compliance documentation), regulator-side tooling (selling to notified bodies who must review AI Act documentation).

Target Customer

Primary persona: a 36-year-old VP Engineering at a Series B AI startup at $14M ARR with EU customers. The compliance team consists of a part-time external counsel. They were quoted $180k by BSI for AI Act readiness work. Will pay $4,500/month Conformant subscription which covers AI Act risk classification + technical documentation + ongoing monitoring with audit-ready artefacts.

Secondary persona: a 47-year-old Chief Compliance Officer at a mid-size enterprise with internal AI deployments (HR automation, customer-service AI, fraud detection). The internal compliance team is overwhelmed by the new AI-specific regulations on top of existing GDPR/SOC2/HIPAA work. Will pay $48,000/year enterprise subscription for the AI-specific layer that complements their existing Vanta or Drata deployment.

Tertiary persona: a 31-year-old founder of a vertical-AI startup at $3M ARR who needs basic AI compliance readiness to satisfy enterprise customer security questionnaires. Will pay $1,800/month early-stage tier for the foundational AI-Act-aligned policy + minimal documentation suite that gets through enterprise procurement.

Product

EU AI Act risk-classification workflow: guided walkthrough of an AI system's intended use, risk categorisation per Annex III (high-risk areas), documented reasoning, supporting evidence repository. Produces the risk-classification artefact in EU AI Act-compliant format.

Technical documentation generator: structured workflow producing Annex IV technical documentation (system description, intended purpose, design specifications, training data summary, accuracy and robustness measures, etc.). Auto-populates from connected sources (model registry, training data logger, monitoring integration) where possible; explicit prompts for human-input sections.

Training-data provenance SDK: lightweight library (Python primarily, with TypeScript planned) instrumenting ML pipelines to log dataset version, source attribution, licensing terms, transformation steps, sample counts, demographic composition where applicable. Produces Article 10 data-governance summaries.

Model card generator: integration with HuggingFace Hub, MLflow, Weights & Biases, AWS SageMaker model registries. Auto-generates and version-controls model cards aligned to Hugging Face model-card spec + AI Act technical documentation requirements.

Post-market monitoring integration: connectors to Helicone, LangSmith, Datadog, OpenTelemetry-based observability for capturing the model-performance and incident data the AI Act requires for high-risk systems. Aggregated reporting per quarter.

Bias-and-fairness testing workflow: structured testing framework for protected-class-attribute bias testing (gender, age, race where applicable, disability), with documented test design + results + remediation tracking. NYC LL144 and Colorado SB 205 compliant outputs.

Incident reporting automation: workflow for capturing serious AI incidents (output causing harm, security breach, model drift causing significant degradation), generating the required regulatory notifications within statutory timelines.

Audit-ready artefact bundle: per-customer downloadable bundle containing all compliance artefacts (risk classifications, technical documentation, model cards, training-data summaries, monitoring reports, bias-test results, incident logs) in formats accepted by EU notified bodies and regulators.

Technical Architecture

Frontend: Next.js + Tailwind for compliance-team-facing dashboard, with deep linking to underlying evidence and audit trail.

Backend: Python + FastAPI on AWS, with strong audit-logging discipline (every action recorded; immutable audit log via cryptographic anchoring).

SDKs: Python primary (instruments ML pipelines), TypeScript secondary (for product-side integration into AI products' UI), with API for languages not covered.

Integrations: HuggingFace Hub, MLflow, Weights & Biases, AWS SageMaker, Vertex AI, OpenAI / Anthropic API usage tracking, Helicone, LangSmith, Datadog, OpenTelemetry.

Document generation: per-jurisdiction document templates (EU AI Act, NYC LL144, Colorado SB 205, ISO/IEC 42001) with structured-data-driven population, PDF + structured JSON exports.

Compliance: SOC2 Type II from year 1 (table-stakes for selling into compliance buyer), ISO 27001 by year 2, GDPR / DPDP compliant data handling.

Business Model & Unit Economics

Three SaaS tiers. Early-stage (■1,800/month): single AI system, foundational risk classification + technical documentation + basic policies; for Series A/B AI startups. Growth (■4,500/month): up to 5 AI systems, full workflow suite, monitoring integration, audit-prep tools; for Series B/C scaling AI companies. Enterprise (■15k-80k/year, custom): unlimited AI systems, multi-team workspace, custom regulator-specific configurations, named compliance advisor, audit-prep engagement included.

Conversion economics: sales cycle 4-9 weeks for SaaS tiers (compliance buyer + technical reviewer + procurement). Conversion from qualified demo: 28%. Distribution: 50% Early-stage, 35% Growth, 15% Enterprise. Monthly churn target under 2.5%.

Gross margin: 76% blended. Major cost lines: infrastructure (~$2/customer/month at scale), compliance-research team for keeping documentation templates current with regulatory changes (~$3.5k/customer/year amortised), customer success and support.

Audit-prep engagements (separately priced): $25k-80k per engagement for structured audit preparation in advance of EU AI Act conformity assessment or US AI audit. ~40% gross margin on engagement labour.

Unit Economics (Year-1 base case)

Year-1 paying subscribers (target)	180 SaaS + 25 enterprise
Blended ARPU	$13,800/year
Year-1 SaaS revenue	$2.4 million
Year-1 engagement revenue	$400k (8 engagements)
Year-1 total revenue	$2.8 million (~■23 crore)
Gross margin	73%
Customer acquisition cost (CAC)	$3,200
Year-1 all-in costs	~■9.5 crore
Year-1 net contribution	~■13.5 crore

Go-to-Market

Channel 1 — Direct outreach to AI-company compliance / engineering leaders (40%): targeted outreach to 500 AI startups with explicit EU presence + 200 mid-size enterprises with internal AI deployments. The market is small enough that named-account targeting works. Conversion target: 60 paying customers in year 1 from outreach.

Channel 2 — Content + thought leadership (25%): substantive content on AI compliance topics (EU AI Act analysis, NYC LL144 implementation guides, ISO/IEC 42001 readiness), conference speaking at AI safety + compliance events, regulatory-update newsletter that builds inbound from compliance buyers.

Channel 3 — Channel partnerships with AI-focused law firms and consultancies (20%): partnerships with DLA Piper's AI practice, BABL.ai, AI Layer (consultancy) for referral arrangements — they sell consulting engagements, we provide the ongoing SaaS layer.

Channel 4 — Vanta / Drata / Secureframe integration partnerships (15%): integration partnerships where Conformant provides the AI-specific layer alongside Vanta's SOC2 layer (Vanta's customer base of AI companies is itself a target segment).

Roadmap (first 12 months)

- Month 1-3: Foundational platform — EU AI Act risk classification + technical documentation generation + basic policy library, SOC2 Type I, first 15 paying customers.
- Month 4-5: Training-data provenance SDK + model-card generator, HuggingFace + MLflow integrations, 50 customers, $400k annualised revenue.
- Month 6-8: Post-market monitoring integration with Helicone + LangSmith + Datadog, bias-and-fairness testing workflow, NYC LL144 + Colorado SB 205 jurisdiction support, 110 customers + 8 enterprise.
- Month 9-10: Audit-ready artefact bundles, audit-prep engagement service launched, ISO/IEC 42001 readiness module, 150 customers + 18 enterprise.
- Month 11-12: 180 SaaS customers + 25 enterprise, $2.8M annualised revenue.

Key Risks

- Vanta / Drata / Secureframe building competitive AI-specific depth — likely 18-30 month threat as the generalist platforms expand into AI compliance. Mitigated by depth advantage (we are AI-specific from day one; they are bolting AI onto generalist platforms) and by integration positioning (we work alongside Vanta rather than competing entirely).
- Regulatory landscape shifting: EU AI Act implementation acts and guidelines continue to evolve through 2026-2027; US federal AI legislation uncertain — both a feature (regulatory complexity drives demand for our product) and a risk (we must stay continuously current). Mitigated by dedicated regulatory-research team (2 senior staff full-time tracking regulations) and by quarterly platform-update cadence.
- Liability if Conformant-generated documentation fails an audit or regulatory review — substantial reputational and possibly legal exposure; mitigated by professional indemnity insurance ($5-10M coverage), by clear scope-of-service contracts disclaiming completeness guarantees in customer's specific factual circumstances, by audit-prep engagement tier providing higher-touch human review for high-risk customers.
- AI Act enforcement slower than projected — actual enforcement timelines have slipped before; if AI Act enforcement remains lax through 2027, willingness-to-pay for compliance tooling depresses. Mitigated by diversification across multiple regulatory frameworks (not just AI Act), by customer-driven demand (enterprise customers requiring AI compliance from vendors regardless of regulator activity), by EU-focused initial market that has clearer enforcement than US.
- AI-specific regulatory frameworks proliferating to point of fragmentation: state-by-state US laws + sector-specific frameworks + international variants could create compliance fragmentation that's hard to address consistently — operational complexity; mitigated by modular platform architecture (per-jurisdiction modules) and by careful customer-onboarding scoping that identifies the specific frameworks each customer requires.