VECTO

Autonomous Intelligence. Human Control.

AI Act Compliance

How Vecto's architecture supports EU AI Act readiness

vecto-os.com · 5 May 2026

Executive Summary

The EU AI Act, enforceable across the European Union, fundamentally changes how organizations must manage AI systems. For any team using AI in software development, project orchestration, or operational workflows, the act introduces obligations around **risk classification, documentation, traceability, human oversight, and post-market monitoring** that did not exist before.

Most existing project management and developer tools were not designed with these obligations in mind. They were built when "AI" meant a chatbot at the bottom of the page — not an agent with authority to write code, plan sprints, or trigger production deployments.

Vecto is built differently. Every agent action is logged with full context. Every decision is traceable to its origin. Every deployment carries a chain of custody from idea to release. This whitepaper explains how Vecto's architecture maps onto the concrete obligations of the EU AI Act, and how organizations can use Vecto as a foundation for their compliance posture.

1. The EU AI Act in Two Pages

What it does

The EU AI Act classifies AI systems by **risk tier** and imposes obligations proportional to that tier:

Tier	Examples	Vecto-relevant obligations
Unacceptable	Social scoring, real-time biometric ID in public	Prohibited
High-risk	AI in critical infrastructure, employment decisions, law enforcement	Conformity assessment · Risk mgmt · Data governance · Documentation · Human oversight · Accuracy, robustness, security · Post-market monitoring
Limited-risk	AI generating content, chatbots, deepfakes	Transparency (label as AI-generated)
Minimal-risk	Spam filters, AI in games	No specific obligations
GPAI (General Purpose)	Foundation models	Technical documentation · Transparency · Compliance with EU copyright

For most organizations using AI for project management, code generation, or operational workflows, the relevant tiers are **High-risk** (when AI decisions affect employment, critical operations, or regulated industries) and **Limited-risk** (general productivity AI).

What it requires (the core obligations)

For high-risk systems, providers and deployers must:

1. **Risk management system** — continuous, documented
2. **Data governance** — quality, representativeness, bias mitigation
3. **Technical documentation** — Annex IV minimum
4. **Record-keeping** — automatic logs of system events
5. **Transparency** — instructions for use, system limits
6. **Human oversight** — design enables meaningful human control
7. **Accuracy, robustness, cybersecurity** — appropriate to risk
8. **Post-market monitoring** — continuous performance + incident tracking

The act takes effect in phases through 2026-2027. **Penalties reach up to €35 million or 7% of global annual turnover**, whichever is higher.

---

2. The Compliance Burden — What's Actually New

Most organizations already have some governance practices. The AI Act adds three things that are genuinely new:

a) Documentation that survives the team

Annex IV requires technical documentation that explains *why* the system was built, *how* it makes decisions, *what* alternatives were considered, and *how* it is monitored. This documentation must remain valid throughout the lifecycle of the system — not just at certification time.

Most tools cannot produce this kind of cumulative, decision-anchored documentation. They produce tickets and commits, not reasoning chains.

b) Human oversight as a design property

The act does not just require human oversight in policy. It requires the system to be **designed** so that human oversight is meaningful. That means agents must operate at a clearly defined

autonomy level, with human intervention points, with the ability to stop, override, or audit.

Most AI tools today operate as black boxes with a single autonomy level. They do not let you tune how much authority the AI has per task or per project.

c) Audit trail as a first-class concern

For high-risk systems, every system event must be logged automatically with sufficient detail to reconstruct what happened. Not "we have logs somewhere." Structured, immutable, exportable, traceable to the actor (human or agent).

Most logging is application-level, not actor-level. It tells you *what* happened, not *who* (or *which agent*) did it, *with what authority*, and *based on which prior decision*.

---

3. Vecto's AI Governance Architecture

Vecto is built on three architectural choices that directly address these new requirements:

3.1 The Living Graph

Every artifact in Vecto — every idea, decision, plan, task, deployment, incident — is a node in a typed graph. Every relationship between artifacts is a typed edge. This means:

- Decisions point to the problems they solved
- Plans point to the incidents that triggered them
- Code changes point to the user stories that motivated them
- Releases point to the security reviews that approved them

The graph is the documentation. Annex IV-style technical documentation can be **derived** from the graph at any point in time — automatically, complete, and traceable.

3.2 The Autonomy Spectrum

Every Vecto agent operates at a configurable autonomy level, from `Observe` (read-only audit) to `Automate` (full autonomy). The level can be set per project, per user, per task type. This gives compliance teams a concrete, measurable answer to "how much human oversight is built into this system?"

Level	Human role	AI authority
Observe	Decision maker	Reports only
Advise	Decision maker	Suggests options
Assist	Reviewer	Acts after approval
Delegate	Boundary-setter	Acts within bounds
Automate	Auditor	Acts independently (rare, by design)

The default is `Assist`. `Automate` is explicitly flagged as exceptional.

3.3 Cryptographically Anchored Audit Trail

Every agent action emits an audit event with: actor identity (human or agent), timestamp, scope (which workspace), action category, sensitivity, severity, and a reference to the decision-context that authorized it. Audit events are immutable, exportable, and signed. **Releases carry HMAC-signed sign-off chains** linking the original idea to the deployed artifact.

This is not an afterthought logging layer. It is a core data structure.

---

4. How Vecto Maps to AI Act Articles

AI Act obligation	Vecto capability
Art. 9 — Risk management system	Risk register as graph nodes; decisions point to risks; periodic risk review built into release flow
Art. 10 — Data governance	Data lineage via graph edges; data sources documented per node
Art. 11 — Technical documentation	Auto-generated from graph; aligned with Annex IV outline
Art. 12 — Record-keeping	Immutable audit log per actor + action; cryptographic integrity
Art. 13 — Transparency	Agent capabilities published per autonomy level; instructions for use built into onboarding
Art. 14 — Human oversight	Autonomy spectrum design; explicit sign-off gates per release; intervention always available
Art. 15 — Accuracy, robustness, cybersecurity	TLS 1.3 transport, AES-256 at rest, RBAC, tenant isolation, signed releases
Art. 26 — Deployer obligations (logging)	All logs exportable; post-market monitoring data produced automatically
Art. 73 — Serious incident reporting	Incident node type with timestamp, severity, and link to triggered decisions

This is not a 100% turnkey AI Act solution. **It is the foundational data architecture on which a compliant program is built.** Combined with appropriate organizational policies, Vecto significantly reduces the documentation and traceability burden.

5. Documentation that Annex IV Asks For (and Vecto Auto-Generates)

Annex IV requires:

1. General description of the AI system
2. Detailed description of elements + dev process

3. Information on monitoring, functioning, control
4. Description of risk management system
5. Description of changes throughout lifecycle
6. List of harmonised standards applied
7. Copy of the EU declaration of conformity
8. Description of the post-market monitoring system

Of these, items **2, 3, 5, and 8** are documentation that traditionally requires manual writing and maintenance. In Vecto, they are **outputs of the graph itself**:

- Item 2 (dev process) → derivable from the chain of IDEAs → EPICs → STORIES → TASKS → RELEASES
- Item 3 (monitoring + control) → derivable from autonomy-level configuration + audit-trail completeness reports
- Item 5 (changes throughout lifecycle) → the graph is, by definition, a chronological record of changes
- Item 8 (post-market monitoring) → derivable from incident-node aggregation + auto-generated KPI dashboards

This is the difference between *writing documentation* and *exporting documentation*.

---

6. The Vecto Compliance Checklist

Use this checklist to evaluate any AI tool against AI Act readiness — including Vecto:

Traceability

- [ ] Every AI action is logged with actor identity
- [ ] Logs are immutable and tamper-evident
- [ ] Logs are exportable in machine-readable format
- [ ] Decisions can be traced from outcome back to original context

Human oversight

- [ ] AI authority is configurable per task / project / user
- [ ] Human intervention points are designed into the workflow
- [ ] AI cannot bypass human review for safety-critical decisions
- [ ] Operators can override or stop AI actions in real time

Documentation

- [ ] Technical documentation can be regenerated at any point
- [ ] Documentation includes the *reasoning* behind decisions, not just outcomes
- [ ] Changes throughout the lifecycle are recorded with context
- [ ] Documentation is not a snapshot — it is a living artifact

Data governance

- [ ] Data sources are documented per AI use case
- [ ] Data lineage from input to output is traceable
- [ ] Bias and quality assessments are recorded

Security

- [ ] Encryption in transit (TLS 1.3+) and at rest (AES-256)
- [ ] Multi-tenant isolation at the database level, not just application level
- [ ] Access control is role-based and least-privilege by default
- [ ] Secrets are encrypted with rotatable keys

Post-market monitoring

- [ ] System events feed an aggregated monitoring view
- [ ] Incidents can be categorized by severity automatically
- [ ] Performance metrics are tracked over time
- [ ] Reports can be generated on demand for regulators

If you cannot check most of these boxes today, you are facing significant work to reach AI Act readiness. Vecto provides this checklist as a foundation, not as the entire solution.

---

7. What Vecto Does Not Do

To be clear about scope:

- Vecto is **not** an AI Act certification. Conformity assessment requires a notified body in many cases.
- Vecto does **not** replace your DPO, your CISO, or your compliance team. It supports them.
- Vecto does **not** automatically classify your AI system risk tier. That requires legal + business judgment.

- Vecto's audit trail is **necessary but not sufficient** for full Annex IV documentation. It is the foundation, not the complete document.

We position Vecto as the **infrastructural backbone** for AI governance. The organizational, legal, and policy work remains yours — but the technical evidence is auto-generated, traceable, and exportable.

---

8. About Vecto

Vecto is a graph-native projectplatform built for organizations that take AI governance seriously. We are headquartered in Belgium, with EU sovereignty as a core architectural choice. Our platform is available as public SaaS, private SaaS, on-prem, and (planned) physical appliance.

Founded in 2026 by a team with backgrounds in cryptography, enterprise security, AI orchestration, and product delivery. Foundation-led governance ensures the platform's mission cannot be diluted or sold off.

For demo, evaluation, or compliance-conversation: **vecto-os.com**

This whitepaper is informational, not legal advice. Consult qualified counsel for AI Act conformity assessment specific to your context.