

AI System Inventory

Vesta Mutual Insurance AS · Register version 1.2 · 29 June 2026 · Owner: AI Governance Lead · Review: quarterly

Worked example for portfolio and training purposes. Vesta Mutual Insurance AS is a fictional company; all data, metrics and names are invented. Prepared by Erik Bernath, Furioso AI Consulting OÜ (furiosoaiconsulting.eu), June 2026. Licensed CC BY 4.0: reuse freely with attribution. This document is informational and is not legal advice.

1. Purpose and scope

This register lists every AI system in use at Vesta Mutual Insurance AS (240 staff; Estonia, Latvia, Lithuania), classifies each under the EU AI Act, and assigns an accountable owner. It is the foundation artifact of Vesta's AI governance programme: the risk assessments, policy controls, monitoring duties and incident procedures all key off this list. Scope covers production systems, embedded AI features in procured software, and sanctioned staff tools. Pilots and experiments enter the register at procurement or before first use with real data, whichever comes first.

2. How the inventory was built

Three discovery passes, repeated quarterly. First, a structured survey of all department heads (what tools make or support decisions in your area, including features inside existing software). Second, a procurement and accounts-payable review for AI-relevant subscriptions. Third, an IT review of SaaS access logs and browser SSO data to surface unsanctioned tools. The third pass matters most: pass one found 6 systems, passes two and three found 3 more, including a transcription tool no manager listed.

Definition used. A system enters the register if it meets the AI Act Art. 3(1) definition (machine-based system that infers from inputs how to generate outputs such as predictions, content, recommendations or decisions). Spreadsheets with fixed formulas and rule-only workflows stay out; statistical and ML models, LLM tools and AI features inside SaaS products go in.

3. The register

#	System	Origin	Function	AI Act classification	Our role	Owner
1	TalentScreen Pro	SaaS vendor (NL)	Ranks and shortlists job applicants from CVs and screening questions	High-risk: Annex III 4(a), employment	Deployer	Head of HR
2	PriceWise v2.3	Built in-house	Risk scoring for term life and health pricing decisions	High-risk: Annex III 5(c), life/health insurance pricing	Provider and deployer	Chief Actuary
3	Aida chatbot	SaaS vendor (DE), LLM-based	Customer service: policy questions, claims status, FNOL intake	Limited risk: Art. 50 transparency duties	Deployer	Head of Customer Ops
4	ClaimSort	Built in-house	Prioritizes incoming claims for handler queues by complexity	Minimal risk (monitored; not an Annex III use)	Provider and deployer	Head of Claims
5	FraudLens	SaaS vendor (UK)	Flags claims with fraud indicators for investigator review	Minimal risk; Annex III 5(b) explicitly excludes financial-fraud detection	Deployer	Head of Claims
6	M365 Copilot	Microsoft	Staff productivity: drafting, summarizing, meeting notes	Minimal risk; Art. 4 literacy and acceptable-use rules apply	Deployer	Head of IT
7	BrandGen suite	SaaS vendor (US)	Marketing copy and image generation	Minimal risk; Art. 50(2)/(4) synthetic-content disclosure applies	Deployer	Head of Marketing

8	Notetaker.ai	SaaS vendor (US)	Meeting transcription and summaries	Minimal risk; GDPR controls dominate	Deployer	Head of IT
9	GitHub Copilot	Microsoft	Code assistance for the 6-person IT team	Minimal risk; secure-coding review rules apply	Deployer	Head of IT

4. Notes on the classification calls

TalentScreen Pro (#1). Annex III 4(a) covers AI used for recruitment and selection, including filtering and evaluating candidates. Vesta is a deployer: Art. 26 obligations apply (use per vendor instructions, trained human oversight, input-data quality, log retention of at least 6 months, monitoring), plus the Art. 26(7) duty to inform workers and their representatives before use, and GDPR Art. 22 discipline because shortlisting decisions significantly affect candidates. A standalone risk assessment exists for this system.

PriceWise (#2). Annex III 5(c) makes AI for risk assessment and pricing of natural persons in life and health insurance high-risk. Because Vesta developed PriceWise internally and uses it under its own name, Vesta carries provider obligations (Art. 16: quality management, Annex IV technical documentation, conformity assessment, registration) on top of deployer duties, and Art. 27 requires a fundamental rights impact assessment for this use. A model card and FRIA exist as companion artifacts.

FraudLens (#5). A deliberate non-high-risk call worth documenting: Annex III 5(b) on creditworthiness explicitly excludes AI used to detect financial fraud. The system stays minimal-risk with standard monitoring, and the exclusion reasoning is recorded here so an authority sees the analysis, not just the conclusion.

Aida (#3) and BrandGen (#7). Both carry Art. 50 transparency duties: customers must be told they are talking to a machine, and synthetic marketing content must be disclosed as AI-generated where required. Neither rises to high-risk.

5. Shadow AI

The SSO review found staff accounts on 2 unsanctioned GenAI tools. Both are blocked at SSO level; the acceptable-use rules in the AI Policy (companion artifact) define the sanctioned alternatives. The register treats shadow AI as a standing discovery category, never as a one-time cleanup: the quarterly IT pass repeats every cycle.

6. Framework mapping and maintenance

- **EU AI Act:** the register implements the classification step that every deployer obligation depends on; classification rationale is recorded per system.
- **ISO/IEC 42001:** serves as the AI asset inventory the management system requires (cl. 4 context, cl. 8 operational planning, Annex A controls on inventory and impact assessment).
- **NIST AI RMF:** implements MAP 1 (context and inventory); owners and review cadence implement GOVERN 1 accountability structures.

Each entry carries owner, last-review and next-review dates in the master spreadsheet version. Any change of vendor, model version, or use purpose triggers reclassification review by the AI Governance Lead before deployment, with material changes going to the quarterly AI Governance Committee.