AI Policy

Vesta Mutual Insurance AS · Policy version 1.0 · Approved by the management board 29 June 2026 · Owner: AI Governance Lead · Review: annual

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 policy sets the rules for selecting, deploying and operating AI systems at Vesta Mutual Insurance AS. It applies to all staff, all departments, and all AI systems in the AI System Inventory (companion artifact), including AI features embedded in procured software. It exists for three reasons: to keep Vesta's use of AI lawful under the EU AI Act and GDPR, to keep decisions that affect people explainable and contestable, and to let staff use productivity AI confidently inside clear lines.

2. Principles

Five rules govern every AI use at Vesta. Lawfulness: no system is deployed outside its legal classification and obligations. Human accountability: a named owner answers for every system, and no decision with material effect on a person is final without human review. Transparency: people are told when AI materially shapes an interaction or decision about them. Proportionality: controls scale with risk, so the pricing model and the meeting transcriber are governed differently. Competence: nobody operates an AI system they have not been trained for (Art. 4).

3. Governance roles

Role	Accountability
Management board sponsor (COO)	Owns the policy, accepts residual risk for high-risk deployments, chairs the quarterly AI Governance Committee
AI Governance Lead	Maintains the inventory and risk assessments, runs classification, coordinates training and incident response, reports quarterly
System owners (named per inventory entry)	Operate the system within this policy, monitor performance, report incidents and changes
Data Protection Officer	GDPR interplay: DPIAs, Art. 22 safeguards, privacy notices, breach overlap in incidents
IT Security	Access control, logging infrastructure, shadow-AI discovery, vendor security review
Works council interface (HR)	Worker information duties for workplace AI (AI Act Art. 26(7)) and consultation where national law requires

4. Classification and intake

Every new AI system, AI feature, or material change to an existing system goes through intake before first use with real data: (1) the requesting owner completes the intake form (purpose, data, affected persons, vendor); (2) the AI Governance Lead screens against the AI Act Art. 5 prohibited-practices list, then classifies (high-risk, transparency-risk, minimal); (3) high-risk classifications trigger a risk assessment, deployer-obligation checklist, and where applicable a FRIA (Art. 27) and DPIA before approval; (4) the system enters the inventory with an owner and review date. The COO approves high-risk deployments; the AI Governance Lead approves the rest.

5. Prohibited uses

Vesta will not deploy systems in the AI Act Art. 5 prohibited categories, and 2 deserve naming because they are marketed to companies like ours: emotion recognition in the workplace (including in interviews or call-center coaching) and social scoring. In addition, company-prohibited regardless of legality: entering client personal data or non-public business data into unsanctioned AI tools; using AI to generate customer communications that are not reviewed by a human; and any fully automated claim denial or policy refusal.

6. Rules for high-risk systems

For every system classified high-risk (currently TalentScreen Pro and PriceWise), the owner operates the full Art. 26 deployer checklist: vendor-instruction conformance, trained human oversight with authority to override, input-data quality control, monitoring with provider feedback, log retention of at least 6 months (Vesta standard: 12), worker information where workplace-relevant, and registration checks. Where Vesta is also the provider (PriceWise, developed in-house), Art. 16 provider obligations apply: quality management, Annex IV technical documentation, conformity assessment and EU database registration, all tracked in the PriceWise compliance plan. The Art. 27 FRIA duty applies to life/health insurance pricing and is on file for PriceWise.

7. Transparency duties

Customers interacting with the Aida chatbot are told at first contact that they are talking to an AI system and can reach a human on request (Art. 50(1)). AI-generated or AI-altered content in marketing is disclosed where Art. 50 requires, and Vesta labels synthetic images as a default beyond the legal minimum. Internal documents drafted with AI assistance carry no disclosure duty, and the author remains fully responsible for content.

8. Generative AI acceptable use

Sanctioned tools: M365 Copilot, BrandGen, Notetaker.ai, GitHub Copilot, each under its inventory entry. Staff may use them for drafting, summarizing, coding and research with three duties: verify before relying (the user owns every output), respect data rules (client personal data only in tools IT has approved for it; never in free public tools), and disclose AI assistance where a recipient would reasonably expect to know. Unsanctioned tools are requested through intake, never adopted directly. The SSO block list is maintained by IT Security.

9. Procurement and vendors

AI procurement adds 6 questions to standard vendor review: What is the system's AI Act classification and the vendor's role? For high-risk: where is the conformity documentation and CE marking? What data trained the system, and what data does it retain from us? How are model updates communicated (Vesta requires 30-day notice and changelogs for high-risk)? What logging do we receive for our oversight duties? What is the exit path (data return, format, deletion)? Contracts for high-risk systems must carry the update-notice and audit-support clauses; the AI Governance Lead reviews before signature.

10. AI literacy (Art. 4)

Every employee receives role-based AI training: general staff (the rules in sections 5 and 8, plus how to recognize and report AI failures), system operators (tool-specific oversight training before access), leadership (obligations, risk acceptance, incident decisions), and IT/security (technical controls). Training happens at onboarding and annually, and completion records with curriculum versions are retained as compliance evidence. Vesta treats the literacy record itself as an audit artifact: who was trained, on what, when, by whom.

11. Incidents, exceptions, review

Suspected AI incidents follow the AI Incident Response Plan (companion artifact) and its severity ladder; every employee can raise one directly with the AI Governance Lead, anonymously if preferred. Exceptions to this policy require written COO approval with an expiry date, recorded in the exception log. The policy is reviewed annually, and within 60 days of any relevant change in law or any SEV-1 incident.

12. Framework mapping

EU AI Act: Art. 4 (s.10), Art. 5 (s.5), Art. 16 and Annex IV (s.6), Art. 26 (s.6), Art. 27 (s.4, s.6), Art. 50 (s.7). ISO/IEC 42001: policy (cl. 5.2), roles (cl. 5.3), operational controls (cl. 8), Annex A controls on acceptable use, impact assessment and supplier management. NIST AI RMF: GOVERN function throughout; intake implements MAP 1.