

What Each AI Tier Actually Decides

Real decisions instead of demos, and the column that keeps getting left off.
Six tiers, not five: a workflow runs the path you wrote. An agent picks its own.

This is not a maturity scale. Organisations run all six at once and add more of them as they mature. Nobody graduates off tier 1.

TIER / WHAT IT DECIDES	EU AI ACT EXPOSURE	CAN YOU UNDO IT?
1 Rules & logic Benefits eligibility · Sanctions screening · Overdraft limits	Often falls outside the Act's definition of an AI system	Yes Read the rule
2 Classical ML Credit scoring · Insurance pricing · Fraud detection	Annex III high-risk. Fraud detection is carved out	Partly Weights inspectable
3 Deep learning CV screening · Imaging triage · Defect rejection	Annex III, or Annex I where embedded in a device	Hard No readable reason
4 Generative AI Claim summaries · Candidate shortlists · Draft denials	Art. 50 transparency, plus Annex III if it decides	Hard Fluent either way
ARCHITECTURE — WRAPS ANY TIER ABOVE		
5 Workflows Routes the claim · Chains score-then-check · Escalates exceptions	Where Art. 14 human oversight can actually be designed in	By design You wrote the path
6 Agentic Sends the denial · Closes the ticket · Moves the money	Inherits the classification of the action it takes	No It chose the path

Simplest to build is not lightest to govern.

Post-Omnibus timing: Annex III stand-alone high-risk from 2 Dec 2027 · embedded and safety-component from 2 Aug 2028 · Art. 50 from 2 Dec 2026.

FOLLOW FREDRIK LINDSTROM FOR MORE

Fredrik Lindstrom

AI and cybersecurity governance for boards, CISOs and senior leaders. 25+ years in enterprise security.

Every promise. Every risk. The truth.