

The claim

A board that authorises an AI system authorises a thing that will not exist for long. The system demonstrated in March — weights, thresholds, training data — is not the system running in June. Retraining, vendor updates and feedback loops rewrite the authorised object weekly, sometimes silently. The question is not whether we approve this system, but what our signature attaches to once it starts moving.

Our position: point-in-time approval of a learning system is void the moment it changes. What a board can honestly authorise is an envelope — the bounds within which the system may change, the monitoring that proves it is still inside them, and the tripwires that revoke authorisation automatically when it is not. A board that cannot produce that document has not authorised a system; it has authorised a memory of one.

The evidence

European law made the envelope statutory. The AI Act defines "substantial modification" as a change not foreseen in the initial conformity assessment that affects compliance or purpose, and requires a fresh assessment for a substantially modified high-risk system. For systems that continue to learn, changes pre-determined at the initial assessment and recorded in the technical documentation are not substantial modifications at all.

EVIDENCE

Under EU AI Act Article 3(23) and Article 43(4), a change to a high-risk AI system that was not pre-determined in the initial conformity assessment and affects compliance or purpose triggers a new conformity assessment; for systems that continue to learn, pre-declared changes recorded in the technical documentation do not constitute substantial modification.

Source: Regulation (EU) 2024/1689, Articles 3(23) and 43(4), text verified 2026 · Geography: European Union · Method: statutory text · SEG-5 · Caveat: applies to high-risk systems; the application timeline of the high-risk obligations was postponed by the Digital Omnibus on AI, Regulation (EU) 2026/1744, in force 27 July 2026.

Read as governance design: the legislator concedes that re-approving a learning system at every change is impossible, and demands instead that the permitted change-space be declared in advance. That is an envelope, in statute. What it leaves open is who watches the boundary.

ISO/IEC 42001:2023, the certifiable AI management-system standard, requires the machinery a moving artifact needs: planned monitoring, performance evaluation, internal audit, management review and continual improvement, for as long as the system operates. ISO/IEC 23894:2023 supplies the risk process inside it. Neither sets a numeric threshold.

EVIDENCE

ISO/IEC 42001:2023 requires organisations operating AI systems to maintain continual monitoring, performance evaluation, internal audit and improvement as ongoing management-system obligations, with ISO/IEC 23894:2023 providing the AI risk-management process across the lifecycle; neither standard fixes numeric drift thresholds.

Source: ISO/IEC 42001:2023 and ISO/IEC 23894:2023 (International Organization for Standardization / IEC), 2023, current 2026 · Geography: international · Method: standard text · SEG-5 · Caveat: management-system standards specify process, not threshold values; certification evidences the machinery, not the safety of any given model change.

The oldest evidence is bank model risk management. SR 11-7 made ongoing monitoring, outcomes analysis and periodic revalidation mandatory for bank models in 2011 — validation as a lifecycle, not an event. On 17 April 2026 it was superseded by the interagency guidance SR 26-2 / OCC Bulletin 2026-13, which carries the discipline to in-scope AI and pointedly excludes generative and agentic AI as too novel to standardise: the supervisors of the world's largest balance sheets refuse point-in-time comfort, and refuse to pretend they have solved the newest models.

The three regimes triangulate the same object: declared change-bounds, running monitoring machinery, supervised revalidation. What none writes — because it is a firm position about board mechanics, not a legal fact — is the third element: revocation tripwires. Bounds without automatic consequences are decoration. The envelope is complete only when pre-agreed conditions — a drift metric out of range, an undeclared retrain, a population shift — suspend the system's authority *without waiting for the next board meeting*.

Evidence cards

CLM-AIACT-PREDETERMINED-CHANGES · SEG-5. Claim: EU AI Act Article 3(23) defines substantial modification as a change not foreseen in the initial conformity assessment affecting compliance or purpose; Article 43(4) requires a new conformity assessment for substantially modified high-risk systems, while pre-determined, documented changes to learning systems are not substantial modifications. Context: Regulation (EU) 2024/1689, text verified 2026. Method: statutory text. Contradictory evidence: scope is high-risk systems; application of the high-risk obligations was postponed by Regulation (EU) 2026/1744 (in force 27 July 2026), so enforcement practice is not yet observable. Causal confidence: none claimed — legal fact. Transferability: EU-regulated deployments; the envelope logic is persuasive elsewhere. Review date: 2026-08-02.

CLM-ISO-AIMS-CONTINUAL-MONITORING · SEG-5. Claim: ISO/IEC 42001:2023 imposes continual monitoring, performance evaluation, audit and improvement as ongoing obligations of an AI management system, with ISO/IEC 23894:2023 supplying the lifecycle risk process; neither fixes numeric thresholds. Context: international standards, published 2023, current 2026. Method: standard text. Contradictory evidence: certification evidences process machinery, not the safety of any specific model change; threshold values remain the organisation's judgment. Causal confidence: none claimed — institutional fact. Transferability: any organisation operating or procuring learning systems. Review date: 2026-08-02.

CLM-MRM-ONGOING-MONITORING-LINEAGE · SEG-5. Claim: US banking supervision has required ongoing monitoring, outcomes analysis and revalidation of models since SR 11-7 (2011), superseded 17 April 2026 by interagency guidance SR 26-2 / OCC Bulletin 2026-13, which extends the discipline to in-scope AI while excluding generative and agentic AI as novel. Context: US supervisory guidance, verified 2026. Method:

supervisory guidance text. Contradictory evidence: scope is US banking organisations; the explicit exclusion means no supervisory template exists for generative or agentic systems. Causal confidence: none claimed — institutional fact. Transferability: model-governed decisions in any regulated balance-sheet business; persuasive beyond banking. Review date: 2026-08-02.

CLM-FIRM-AUTHORISATION-ENVELOPE · SEG-1 (firm position). Claim: point-in-time approval is void for systems that learn; board authorisation must attach to an envelope of declared bounds, monitoring obligations and automatic revocation tripwires, and is incomplete without all three. Context: firm operating practice in the corridor. Method: interpretation built on the three cited regimes, not measurement; no client governance outcomes cited. Contradictory evidence: for static models under change-freeze, point-in-time approval with re-approval at each release can be honest; tripwire design is unvalidated as a standard instrument. Causal confidence: none claimed. Transferability: bounded — strongest for cross-border groups deploying learning systems at operating distance. Review date: 2026-08-02.

The limits

The board sits in Paris or Geneva; the learning system runs in Abidjan, scoring credit or screening fraud. It will never inspect the system and governs only what crosses on paper. If that is a quarterly slide saying "the AI is performing well," it has authorised an artifact it cannot see. What must cross is the envelope report: every retrain, threshold move and vendor update marked inside or outside the declared bounds; drift metrics against their limits, each with an owner and a cadence; boundary events and what they triggered; and who in Abidjan can suspend the system tonight without a call to Paris. The regulatory floor is asymmetric: the AI Act binds the European holding, and no equivalent statute operates in the WAEMU jurisdictions where the system runs, so the envelope exists only if the group imposes it by governance and contract, including on vendors who retrain upstream models without asking. Drift is also faster in the corridor: informal-market data, currency events and thin credit histories move the input distribution more violently than vendor defaults assume, so a cadence copied from a European plan understates corridor drift by design.

Sources and limitations

Sources and limitations. The institutional facts rest on Regulation (EU) 2024/1689 Articles 3(23) and 43(4), with the application timeline amended by Regulation (EU) 2026/1744, on ISO/IEC 42001:2023 and ISO/IEC 23894:2023, and on the supervisory lineage from SR 11-7 to SR 26-2 / OCC Bulletin 2026-13 — each carried with its caveat above. The envelope doctrine — bounds, monitoring, tripwires as the object of board authorisation — is the firm's position, graded as interpretation: the three regimes motivate it, none mandates it in those terms, and we cite no client outcomes because none has passed our evidence-release gate. The transferability boundary is explicit: strongest where a learning system runs at operating distance from the board that authorised it; weakest for frozen models under disciplined release control. This note is not valid as legal advice on conformity assessment, nor as a claim that any certification or supervisory regime, by itself, makes a learning system safe to authorise.

What to do on Monday

1. The company secretary: check whether the last AI authorisation names an artifact or an envelope. If it names only a version demonstrated at a past meeting, it expired at the first retrain.

2. The chief risk officer: write the tripwires and name the officer in the operating company who can pull them without convening the board. Conditions, maximum delay, local authority: a revocation that waits for a board calendar is not a control.

3. The model owner: produce the boundary log since the last meeting. Each retrain, threshold move, population shift and vendor update, marked declared or discovered. Zero boundary events deserves the suspicion an audit with zero findings deserves.

4. The procurement lead: check whether the vendor contract obliges disclosure of model changes against our envelope, in days rather than quarters. If not, the envelope has a hole the size of the vendor.