Give AI reviewers the authority to hold a decision

Adrian Dunkley | 2026-11-15 | Related Council publication · model briefing

Print edition of the public web text, prepared 20 November 2026. Open the source publication.

prepared by AI Boss Limited

A risk committee can test approval timing, shared failures and the limits on delegated execution before widening an AI pilot.

By

Define when the workflow must wait

Ask each process owner to name the first action that is costly to reverse. Require an authorised reviewer to inspect the assumptions behind that action. Set a review deadline and keep the action on hold if the reviewer is unavailable. Track completed reviews before execution separately from checks made after the fact.

Stress several controls together

In Adrian Dunkley’s illustrative analysis, expert supervision reduced expected annual loss to US$317,000. Halving staff time and the share of reviews that checked assumptions material to approval, while doubling the chance of a shared failure, raised expected loss to US$447,000. The human baseline was US$431,000. The threshold exceeded in about 1% of simulated years rose from US$1.42 million in the human workflow to US$1.49 million under combined stress.

In a test environment, simulate a failed supplier feed or approval service. Record the number of affected actions, time to stop, and whether the rollback succeeds. Ask the accountable owner to sign off on any remaining risk before raising transaction limits. Read the model summary and stress-test results.

CAIRMC is a council of the Caribbean AI Association.Browse more research · Browse by topic