

AI governance · Decision review

When Should We Revisit an AI Decision?

A review date can be a real control. It stops being one when nothing is attached to it, and when the events that change the answer arrive between the dates with nobody watching.

What you finish with: A scheduled review with its evidence and its accountable owner named in advance, and three or four written event triggers for one AI decision you have already funded, each carrying the event, the role who sees it first, what they bring, and whether it reopens, is logged, or escalates immediately.

[Open this research online](#)

DECISION REVIEW / TWO ENTRY POINTS

A date and a change can both trigger review.

ON THE CALENDAR

Scheduled review

An agreed date, named owner and current evidence.

WHEN FACTS CHANGE

Event trigger

A defined change, an observer and the action to take.

↓ ↓
Triage → record evidence → decide

Escalate a new error class or control failure immediately. Keep the current control until its authorized owner changes it.

Take an AI decision your organization has already made. It probably carries a review date: twelve months out, at renewal, or annually in the governance pack.

A scheduled review is a control when it does two things: examine named evidence, and produce a decision or an action. Set that way, with the evidence named in advance and someone accountable for bringing it, the date earns its place. A calendar entry with nothing attached to it does not. It produces a meeting that describes what happened rather than a decision about what to do next.

The second problem is timing. Some things that change an AI decision are schedulable and belong on the calendar, renewal dates and planned role changes among them. Others arrive when they arrive: a price change mid-term, a platform release that absorbs what you bought, a volume shift that breaks the review design. Waiting for the next date to notice those is a choice, and it is worth making deliberately rather than by default.

So use both. Keep the scheduled review and give it evidence and an owner. Add a small number of event triggers for the things that will not wait. Four triggers on one page is a working control; a longer list postpones the harder question of which events actually matter here.

01 What tends to reopen a decision

These are candidates to choose from, not a list to complete. Pick the few that fit the decision in front of you.

The commercial terms move. A price rises, a consumption charge appears, a bundle is unbundled, or renewal terms change. Renewal itself is a date, so put it on the calendar. What is not a date is a change inside the term, and that is the part a trigger covers.

The platform absorbs the capability. Your existing platform now does natively what you bought or built separately. Paying for both is a position nobody chose.

The volume or the mix changes. The arithmetic that justified the design was done at one volume. At a materially higher volume the review the design depends on may no longer fit the window it has to fit inside.

The measure moves, or fails to. The improvement you committed to did not arrive by the checkpoint you set, or it arrived and nobody can explain it. Both are worth examining.

A new error class appears. Not more errors, a different kind: an output wrong in a way the review point was not designed to catch. This one escalates immediately. It is not a candidate for a threshold that has to be sustained before anyone acts.

The accountable owner changes. A planned departure is schedulable. An unplanned one is not, and the first move is interim accountability rather than a search.

02 What a trigger has to contain to work

"If costs increase significantly, we will review" is a sentence, not a control. Nobody knows who is watching, what counts as significant, or what happens next. Four elements fix that, and all four fit on one line.

The event, stated so two people would agree it had happened. A threshold, a named change, a date. "Unit cost per processed item rises above X" rather than "costs increase."

Who sees it first, by role. Not necessarily the decision owner. It is whoever reads the invoice, runs the queue, or receives the release note. Their job is to raise it, not to judge it.

What they bring. One or two artifacts, named in advance, so the reopening does not start with a fortnight of assembly.

What happens next, and how fast. Reopened for a decision, logged and carried to the scheduled review, or escalated immediately. Anything describing a control failure or a new error class takes the third route regardless of how the other triggers are written.

03 Example decision record: an accounts payable module

A finance function enabled exception routing and cost code suggestion inside an accounts payable module it already licensed, tested it for one cycle, and kept it. The decision record set a scheduled review at twelve months, with three items of evidence named for that meeting, and four event triggers for what would not wait.

Trigger	Who sees it first	What they bring	What happens next
A consumption or activation charge appears on the module invoice, or terms change inside the contract term	Accounts payable manager, who approves the invoice	The invoice line and the prior period's for comparison	Reopened. Cost was the basis of the original approval
Suggested codes corrected by clerks exceed the rate recorded in the test cycle	Accounts payable team lead	The correction count and the month's exception volume	Reopened if sustained over two months. A correction of a kind not seen in the test cycle escalates immediately, without waiting for the second month
Monthly exception volume rises above 450, from the 300 the review burden was sized on	Accounts payable team lead	The volume trend and current review minutes per item	Reopened. The design was sized at 300 and the close window has not grown
The accountable owner changes role	Finance director	Nothing. The trigger is the vacancy	The finance director holds interim accountability from the vacancy; permanent reassignment within four weeks

What fired. In month three the vendor introduced a per-suggestion consumption charge inside the contract term, below the level anyone had thought to watch for. The accounts payable manager

approving the invoice noticed, because she had a line telling her to. In month five volume reached 470 and the review time per cycle passed the window. Both required attention before the scheduled review.

What the volume trigger did not do. It did not authorize a change of control. Reviewing every suggestion was the control the original approval rested on, so the team kept it while the owner worked the three options that were actually available: reduce the scope of what the feature handles this cycle, add reviewer capacity from within the close team, or design and validate a proportionate alternative such as a stated sample with defined criteria.

The third route needs evidence and an approval before it operates, not after. Moving unreviewed items to a post-posting check because the queue got long is a control change made by circumstance, and it is the kind of drift the trigger exists to prevent.

What the measure did. The improvement committed to at approval was total clerk hours on the exception workload below the measured baseline by the month six checkpoint. At the month six checkpoint the hours were unchanged from baseline. That is a missed target at a named checkpoint, with the target and the checkpoint written down at approval, which is why it is worth examining rather than a general anxiety about a steady number.

04 What monitoring should cost

Use existing routines where they provide adequate coverage: an invoice already being approved, a queue already being run, a report already being read. If they do not provide it, compare the monitoring cost with the exposure and decide whether to fund the monitoring, reduce the exposure, or stop the activity. Record that decision with the accountable owner.

That is a real decision with a real answer, and it is better made once, deliberately, than left to whoever notices the gap later.

05 What to do now

Take one AI decision your organization has already made and funded. Give its scheduled review the two things that make it a control: the evidence to be examined, named in advance, and the person accountable for bringing it. Then write a small number of event triggers, three or four, for the changes that will not wait for that date, each with the event stated so two people would agree it had happened, the role who sees it first, what they bring, and whether it reopens, is logged, or escalates immediately.

Give the list to the people named. A trigger nobody has been told about is not a control.

If you cannot name who would see the first event, you have found something more useful than a trigger list. The decision has no operational owner, and that is the thing to fix first.

Sources. Built from M. Pact material: the re-evaluation trigger and decision-record fields in MPA-RN-2026-001 and the M. Pact Decision Record Schema, the review timing and effort calculations in the Process Boundary Guide, and the ownership and evidence tests in MPA-RN-2026-002.

About this research

© 2026 M. Pact Advisory Services, LLC. All rights reserved.

This material is proprietary to M. Pact Advisory Services, LLC and is provided for executive advisory, strategic decision support, and internal planning purposes. It reflects practitioner judgment informed by experience across enterprise strategy, finance, technology, and operating models.

The frameworks and perspectives in this research are M. Pact Advisory's independent analysis and professional judgment at the time of publication. They are intended to inform executive thinking and decision-making, not to prescribe specific actions for any individual organization.

This research does not constitute legal, accounting, investment, or regulatory advice and should not be relied upon as such. Consult appropriate professional advisors before making decisions based on the information contained here.

The analysis is developed independently by M. Pact Advisory and is not influenced by third parties. M. Pact Advisory accepts no vendor compensation for placement, coverage, or rating. References to market practices, categories, or external concepts are illustrative and do not imply endorsement or affiliation.

This material may not be reproduced, distributed, or transmitted, in whole or in part, without the prior written consent of M. Pact Advisory Services, LLC.

This research may not be used as input for the training, fine-tuning, or development of generative artificial intelligence systems, machine learning models, algorithms, or related technologies without prior written permission.