The claim

Every AI initiative is born owned by enthusiasm. A sponsor finds budget, a pilot is framed as learning, and for a quarter or two that framing is honest: small spend, unknown outcome, costless exit. The decision arrives later, and rarely reaches the agenda in these words: who owns this money now that it recurs? The subscription renewed. The integration has a maintenance line. Two people spend most of their week on prompts, review and data plumbing. The experiment has acquired the one property experiments are not supposed to have — permanence.

Our position: AI spend moves under the CFO's discipline the moment it recurs in the cost base. Not when the technology matures, not when the organisation feels ready — when it recurs. From that moment it deserves what every other recurring commitment gets: a business case with a named benefits line, an amortisation view of anything capitalised, benefits tracking against the case, and kill criteria exercisable at renewal. Treating recurring AI spend as perpetual innovation budget is not open-mindedness; it is how pilots become permanent costs, because innovation budgets are judged on activity and learning, and no line judged on learning ever fails a renewal.

The evidence

The standard-setters are further ahead than most operating models. The IFRS Interpretations Committee concluded in March 2019 that a customer's right to access supplier software hosted in the cloud is typically a service contract, not a software intangible. In April 2021 it finalised the harder question — the configuration and customisation costs poured into making a SaaS platform fit the business — and the answer is symmetrical.

> EVIDENCE
>
> The IFRS Interpretations Committee concluded that cloud software access is typically a service contract rather than an intangible asset (March 2019), and that configuration and customisation costs in a SaaS arrangement are generally expensed unless they create an intangible asset the customer controls under IAS 38 (April 2021).
>
> Source: IFRIC agenda decisions, March 2019 and April 2021, under IAS 38; IFRS Foundation · Geography: IFRS jurisdictions · Method: authoritative interpretation of accounting standards · SEG-5 · Caveat: agenda decisions are explanatory material applied to entity-specific facts; some elements of AI spend — controlled models, owned code, certain licences — can still qualify for capitalisation under IAS 38.

US GAAP reaches a similar structure by another route.

EVIDENCE

Under FASB ASU 2018-15 (ASC 350-40), implementation costs in a cloud computing arrangement that is a service contract follow the internal-use software stages: application-development-stage costs are capitalised and amortised over the arrangement term; preliminary and post-implementation costs are expensed.

Source: FASB Accounting Standards Update 2018-15, issued August 2018 · Geography: United States (US GAAP) · Method: authoritative accounting standard · SEG-5 · Caveat: US GAAP treatment differs from IFRS in what is capitalised; classification depends on entity-specific facts and the arrangement's terms.

The detail matters less than the shared message: this spend is ordinary. It is classifiable, schedulable, amortisable. Nothing in the accounting literature recognises a category called innovation that suspends the discipline.

The oversight side has moved too. The Center for Audit Quality's July 2024 resource places generative AI squarely inside the financial-reporting perimeter, with question sets for audit committees on governance, data, deployment, monitoring, fraud and the regulatory environment. The profession's oversight bodies already treat AI spend and AI use as auditable, questionable, ordinary — a subject for the audit committee's standing machinery, not a protected species.

What the standards and the guidance do not supply is the trigger. That is our judgment, and we state it as a test: recurrence. The second renewal of a subscription, the first budget cycle in which the spend appears without being re-argued, the first headcount whose job description assumes the tool exists. Any of these ends the experiment, whatever the deck says. The measurable symptom of failure is a cost line that grows year on year while its benefits case is still described as learning.

In the corridor the ownership question has a geography. The AI stack of an Abidjan or Dakar operating company is typically contracted by the holding: group licences negotiated in Paris or Geneva, priced in euros or dollars and allocated down through recharges. The local CFO sees an allocation, not a decision. Kill criteria, where they exist, are exercisable at a renewal negotiated two time zones away, for a bundle in which the opco's usage is a rounding line. The costs are real in local currency; the benefits are supposed to appear in local operations that nobody instrumented.

Evidence cards

CLM-IFRIC-SAAS-EXPENSE · SEG-5. Claim: under IFRS, cloud software access is typically a service contract rather than an intangible asset, and SaaS configuration and customisation costs are generally expensed unless they create an intangible asset the customer controls under IAS 38. Context: IFRIC agenda decisions of March 2019 and April 2021. Method: authoritative interpretation of accounting standards. Contradictory evidence: agenda decisions apply to entity-specific facts; some AI-related expenditure (controlled models, owned code, certain licences) can qualify for capitalisation. Causal confidence: none claimed — institutional fact. Transferability: IFRS jurisdictions; the classification logic is persuasive elsewhere. Review date: 2026-08-02.

CLM-USGAAP-CCA-STAGES · SEG-5. Claim: FASB ASU 2018-15 requires implementation costs in a cloud computing arrangement that is a service contract to follow internal-use software stages — application-development costs capitalised and amortised over the arrangement term, preliminary and post-implementation costs expensed. Context: ASU 2018-15, issued August 2018, ASC 350-40. Method: authoritative accounting standard. Contradictory evidence: treatment differs from IFRS; classification is fact-dependent. Causal confidence: none claimed — institutional fact. Transferability: US GAAP reporters and their subsidiaries. Review date: 2026-08-02.

CLM-CAQ-AUDIT-COMMITTEE-GENAI · SEG-4. Claim: the Center for Audit Quality's July 2024 resource places generative AI within audit-committee oversight of financial reporting and internal control, with question sets on governance, data, deployment, monitoring, fraud and the regulatory environment. Context: CAQ, "Audit Committee Oversight in the Age of Generative AI", July 2024. Method: professional-body practice aid. Contradictory evidence: a practice aid, not an authoritative standard; scoped to financial reporting rather than all AI spend. Causal confidence: none claimed — institutional fact. Transferability: any audit committee; strongest for US-listed structures. Review date: 2026-08-02.

CLM-FIRM-CFO-RECURRENCE-TEST · SEG-1 (firm position). Claim: AI spend moves under CFO discipline — business case with a named benefits line, amortisation view, benefits tracking, kill criteria exercisable at renewal — the moment it recurs in the cost base; continued innovation-budget treatment past recurrence converts pilots into permanent costs. Context: firm operating practice in the corridor. Method: interpretation, not measurement; no client cost data cited. Contradictory evidence: genuinely exploratory research spend can legitimately recur under an R&D governance regime with its own gates; the test targets operating deployments, not research. Causal confidence: none claimed. Transferability: bounded — strongest for holding structures with recharged group contracts. Review date: 2026-08-02.

The limits

The discipline has to be built twice, and the second build is where it fails. A group business case proves nothing about Abidjan, so a local benefits line and a local baseline are needed or nothing is tracked. Local statements are drawn under a different accounting framework than the group's IFRS reporting, so the same spend can look like an asset in one book and an expense in the other. A holding board that has never asked how the group's AI recharges land in the opco's accounts is not overseeing AI spend; it is distributing it.

Sources and limitations. The accounting facts rest on the IFRIC agenda decisions of March 2019 and April 2021 under IAS 38 and on FASB ASU 2018-15, both verified against the standard-setters' texts on 2026-08-02; the oversight fact rests on the CAQ's July 2024 publication, a practice aid and graded accordingly. No consultancy material is cited. The recurrence test is the firm's position, graded as interpretation: it is not a measured result, and no client cost or benefits data appear here because none has passed our evidence-release gate. Boundaries are explicit: this note is not accounting advice on any specific arrangement — capitalisation outcomes are fact-dependent and belong with the entity's auditors; the corridor observations on group recharges are operating reasoning, not statistics; and the position does not argue that research spend must show quarterly returns, only that recurring operating spend must stop hiding behind the word innovation.

What to do on Monday

1. The group financial controller lists every AI line appearing for a second budget cycle and names who approved its change of status. If a line moved from innovation to run-rate without a decision, the move happened by default, and default is not a governance mechanism.

2. The CFO reopens the balance-sheet treatment of each recurring line: what was capitalised, under which standard, over what life, and what was expensed. The test is whether the treatment reflects the applicable interpretation rather than the sponsor's preference for a smaller P&L hit.

3. Each sponsor of a recurring line delivers a benefits schedule naming one reporting line, a baseline and a date. Where the answer is a narrative about capability building, apply the recurrence test and fund it, if at all, as a dated experiment.

4. The contract owner writes kill criteria into the next renewal and confirms that the entity bearing the cost can exercise them. A kill right held where the cost is not felt will not be exercised.