

Where Fiduciary AI Actually Dies

It isn't the IT department. It's the examination room.

Ethan Mollick argued in *The Economist* that corporate AI suffocates inside the IT department, where a mandate to minimize risk leaves no room for the experimentation AI demands. He is right about corporates. In regulated fiduciary, the death is quieter and more expensive.

AI doesn't die in IT, it dies two years later when a beneficiary's adviser asks why a discretionary distribution was made and no one can reconstruct the reasoning. The deliberation existed. The record didn't. That gap is not an IT problem. It's a fiduciary one.

Fiduciary deliberation is the central act of the profession. It is, by nature, contextual, nuanced, human. What AI can now do is document that deliberation in real time: structure the evidentiary trail, connect each decision to its parameters, the alternatives weighed, the justification applied. The judgment stays human. The record becomes unassailable.

Most trust companies haven't asked the right question yet. They ask whether AI can accelerate KYC, whether it can cut onboarding time. Those are operational efficiency questions. They were the right questions in 2023.

In 2026, the question that matters is different: what is the contemporaneous, defensible, traceable evidence of fiduciary judgment exercised?

That question is not rhetorical, and the ground is already shifting. CRS 2.0 took effect in Switzerland on 1 January 2026, with the first reporting due in 2027 — raising the granularity of what gets examined and the due-diligence standard expected of reporting institutions. But the transparency regime is the ambient pressure, not the sharp edge. The sharp edge is older and quieter: when a discretionary decision is challenged years later — by a beneficiary's adviser, or under the proper law of the trust — a trustee's protection turns not on the outcome, but on contemporaneous evidence that the relevant considerations were actually weighed at the time. The reasoning either exists in the record, or it doesn't. Structures are increasingly tested not on returns, but on decision architecture.

The firms that have built that architecture, where every fiduciary decision automatically generates its documentation, are already in a different position: defensible at any point in time, by any examiner, on any decision.

The firms that haven't built it are carrying a liability they haven't priced yet.

Frédéric Sanz

Founder & CEO, FiduciaCorp
20+ years UHNW fiduciary management · Switzerland

FIDUCIACORP.COM