

Decision Accounting

Your organization has begun delegating decisions to software. It has not begun accounting for them.

The position you are in

AI agents in your business now price, approve, remediate, procure and communicate — acts that until recently required a person with a title and a signing limit. Your controls have guardrails, built for software. As agents are granted more decision rights, those guardrails must grow in alignment — and today most stop at what an agent may touch rather than what it may decide. Your systems show which service account touched which record. They cannot show what the agent decided, under whose authority, who approved it, or what it cost. You remain accountable for all four.

Why your existing controls do not close this

- **Permissions are not authority** — An entitlement defines what an agent is able to touch. It says nothing about what the agent chose within that space — and almost every consequential thing an agent does is a permitted action.
- **Model validation governs the model, not the decision** — Validation governs an artifact at rest. Authorization is a property of the action, not the weights. In April, U.S. regulators replaced SR 11-7 and placed agentic AI outside its scope, directing institutions to a broader control environment not designed to control the actions of agents.
- **Operational dashboards are not denominated in money** — Latency, containment and error rates describe how an agent is running, not what its running is worth. An agent can score well on every operational measure while destroying value: throughput up 40%, error rate near zero, dashboard green for nine days — because it was buying volume with margin.

The idea, in one paragraph

When Pacioli codified double-entry bookkeeping in 1494, the innovation was not arithmetic. It was accountability: every economic event recorded, every entry traceable, every steward of the firm's money answerable for it. Enterprises are now custodians of a second scarce resource — delegated decision-making authority — and they hand it to synthetic actors at scale without a ledger. Decision accounting extends the discipline: every consequential agent action recorded, expressed in the language of the business, checked against an explicit mandate owned by a named human, and closed against what it produced.

Where does your organization sit?

Level		What you can honestly say today
L0	Unobserved	“Agents act; we cannot reconstruct what they did.”
L1	Logged	“Vendor logs exist, in system vocabulary. We can trace calls; we cannot read decisions.”
L2	Reconciled	“We know what our agents did, in the language of our business.”
L3	Governed	“Every agent acts under an explicit mandate; breaches surface with a number attached.”
L4	Attested	“We can evidence to our board and regulator what our agents decided, under whose authority, and with what consequence.”

Most large institutions sit between L1 and L2 — well controlled at the infrastructure layer, unable to render agent activity in business terms. We expect supervisory expectation to converge on L3 within two to three years.

Five questions worth asking this quarter

1. Which AI agents can write to systems that feed our numbers, and who owns each one by name?
2. What is each one authorized to decide — in metrics and money, not in system permissions — and who approves beyond that?
3. When they act, what record do we hold, and is it in our vocabulary or our vendors'?
4. Are the delegations working — better or worse than what they replaced?
5. How concentrated are we — how much of our agent workforce depends on one model or one vendor?

If those cannot be answered today, the gap is not a technology problem. It is an accounting one, and it belongs on the same agenda as any other delegation of authority.