

What Observable Governance delivers in healthcare and life sciences

Benefits, growth and savings. Continuum as the path. Written in the voice of a working AI programme note.

GOVERNANCE FLOWS DOWN · OBSERVABLE FLOWS UP

Start here

Hospitals, payers and life-science teams already have chat tools and first agents. People already use unofficial ones — on notes, protocols, safety mail and trial files. Supervisors and data-protection officers now ask two things: who may touch the patient or trial file, and how do you know?

Use is common. Profit and live clinical or operational impact are not. McKinsey finds 88 in 100 firms use AI. Only 37 in 100 can show any effect on operating profit. About 6 in 100 see a material effect.

Life sciences sits in McKinsey's high-impact set for generative AI. Most of that value is still in trials. The official tool cannot be halted. It cannot connect to the record. Nobody can look up last month's yes or no. So the Caldicott guardian, the DPO or quality will not bless it. Staff keep the unofficial tool.

Health data has extra locks. GDPR is the floor. France requires a certified health host (HDS) and, for public cloud, often SecNumCloud. Germany limits cloud processing of health and social data and asks for BSI C5. The European Health Data Space expects secondary-use health data to stay in the Union. An "EU region" on a US-operated stack often fails those tests.

What you get if you fix that: the assistant is allowed to finish the job on a named path. Growth is the process you can put it on — coding a trial, coding a claim, drafting a letter, triaging a referral. Saving is the waiting and the duplicate tool you stop running.

This note does not invent "saves \$X per hospital." McKinsey does not publish that. Size the first process. Then copy it.

Words used here

- Observable Governance — the rule that applies to a live step can be shown, and last month's allow or deny can be looked up, with the rule version.
- Continuum — the shared path: owner, rules, door into the official system, record.
- Record — EPR, claims system, safety database, trial master file, LIMS.
- Health host rules — HDS in France, C5 and the social-code cloud rule in Germany, EHDS for secondary use.

1. What it delivers — three tests, in a hospital or a lab




Figure. Policy down onto the note, the claim, the trial step. Proof back up to quality, the DPO and the board.

Test	What you can show
Authority is executable	A named stop owner can halt an agent on a record, a safety file or a trial path. The rule exists as intent and as code, same version.
The estate is bound	The assistant uses the same named door as every other caller — API and, for agents, a governed MCP. Unscoped tools fail closed. If the rules engine is silent, the step is refused.
The loop closes	Allow or deny writes who (clinician, coder, bot or agent), what, which rule version, which door. The next rule cites those records. The quality pack is that look-up.

2. Benefits by the people who feel them

Business line — growth they can ship

- The assistant can update the real record, not only draft in a side window.
- The unofficial tool becomes a named tool — or it stops.
- Quality and the DPO can answer from a look-up, not from a workshop.

Risk, compliance, audit — a pack that exists

- One record serves the AI Act, health-host rules and the board.
- You can show last month's updates stayed in the residency tier you promised.
- “Just this once” has an end date.

Operations and technology — one plane

- No second “AI safety office” beside IG and quality. One path.
- If the host or the model must change, the work has a named exit — not a war-room.

Customer and colleague

- They are told when they are talking to an AI (Article 50, live from 2 August 2026).

- A wrong step can be found. The record is searchable.

3. Growth opportunities — where the path pays

McKinsey puts life sciences among the sectors with the largest generative-AI share of revenue if uses are live. Hospitals and payers see the same pattern as every other industry: time saved at the desk, little in the official process until the tool can touch the record. That is a pool, not a promise.

Domain	What Observable Governance unblocks
Clinical and operational admin	Referral, discharge and inbox agents that may read and, where the rule allows, write the record — not a chatbot that stops at a summary.
Access and medical affairs	Field and medical assistants that can be shown to compliance. Growth without a mystery model on a doctor's laptop.
Coding, claims and prior auth	A coding or triage agent on a named door. The file stays on the health-host tier. The deny is visible.
Safety, quality and trials	Safety-signal and trial assistants on a halt-able path. High-risk Annex III (work, essential services) clocks from 2 December 2027 — build the path now.
Engineering	Build agents that cannot land unsigned artefacts into a validated system. Throughput without a shadow pipeline.

Growth rule: pick the process that is already leaking into unofficial tools. That is the demand signal. Do not start with a platform programme.

4. Savings — what you can defend, and what you cannot

What the research lets you say

- Waiting. On programmes McKinsey watched, 30 to 50 percent of build time was permission and rework. A shared path cut that wait. They have seen approval time drop by as much as 90 percent.
- Missing profit. 88 in 100 use AI. 37 in 100 show any profit effect. About 6 in 100 show a material one.
- When it is live. A selected McKinsey client group that changed how work runs saw about \$3 back per \$1 after the tool was live. Not a promise for a demo.
- The pool. McKinsey's sector work puts a large share of generative-AI value in life sciences and in operations. A live path takes a slice. A stranded trial takes none.

Savings you can size on a first process

- Hours not spent rebuilding the same risk pack for each pilot.
- Unofficial licences, plugins and duplicate prompt-spend retired once the attested assistant does the job.
- Fewer dual-running teams (official bot plus the private one that actually works).
- Incident and audit cost: reconstruct-from-Slack versus query-the-store.
- Host and model exit drills that are paper today and a switch tomorrow (HDS / C5 / EHDS).

What this note will not invent

There is no honest industry-wide “saves 12 percent of opex” figure. Size the first process: hours in the permission queue, unofficial tools to retire, letters or codes written back per week once live.

5. How the saving and the growth show up in 90 days

- Day 0 — stop owner named on one process (inbox, coding, safety, a trial arm).
- Days 1–30 — unofficial tools on that path listed. Dual-form rule. Residency and key tier set.

- Days 31–60 — attested door live (API + governed MCP). Unscoped tool-call fails closed.
- Days 61–90 — query works. First rule change cites live records. Health-host file and AI Act transparency file point at the same store. Unofficial twin retired or denied.

That is the first saving (permission time and a retired twin) and the first growth (write-back on a live record). Clone the path. Do not buy a second tower.

6. Decision

Put this under the next agent that wants to touch a patient, claim or trial file.

Fund one path. Not a new framework.

Growth is the process you are then allowed to scale. Saving is the queue and the unofficial twin you stop running.

If you cannot look up last month's yes or no on a record or a safety file, you are still in the group that uses AI and cannot show it in the official process.

— KanTec Platform Engineering · Healthcare and life sciences brief —