AI COMMITTEES

AI committees fail on throughput, not judgment.

Why governance backlogs are arithmetic, and what actually fixes them.

For CIOs, Chief AI Officers, and the people who chair these meetings.

START WITH THE ARITHMETIC

Do the math on your own committee

36

Decisions per year — and that's the optimistic number. Twelve meetings a year, three tools discussed with any real depth per meeting. Headcount doesn't change this; reviewer attention does.

Most committee chairs I talk to are honest that two tools is closer to the truth, and that two of the twelve meetings don't happen. Which puts the real figure nearer twenty.

Now compare it to your intake.

THE GAP

Capacity grows linearly. Intake doesn't.

Three-quarters of US health systems use at least one AI application; half run three or more (Eliciting Insights, n=120). 93% deploy third-party AI (KLAS/CCM, n=27). Physician use went from 38% to 81% in three years (AMA, n=1,692).

Meanwhile the committee gained one meeting a quarter and two new members.

THE UNCOMFORTABLE VERSION

Meeting twice as often does not double throughput, because the constraint is reviewer attention, not calendar slots. In the committees I have watched, the gain is marginal and intake outruns it inside a year.

WHERE THE TIME ACTUALLY GOES

Almost none of it is judgment

→ **Chasing evidence that was never requested clearly**

The single largest time sink. A vendor sends what they have; the committee needs something else; three weeks pass.

→ **Re-explaining the tool to people who weren't at the last meeting**

A committee of 12 with real jobs has different attendance every month.

→ **Relitigating decisions the committee already made**

“Do we accept a missing subgroup analysis?” — asked again, from scratch, every time it comes up.

→ **Discussing tools that never needed a discussion**

Scheduling optimizers and templating tools consuming the same slot as an autonomous clinical model.

The actual deliberation — the part that requires twelve senior people in a room — is a small fraction of the meeting.

THE FIX

Stop routing everything through one path

RULE-ROUTED · DAYS

No patient-facing output, no PHI leaving the boundary, advisory impact, human in the loop.

Templating, scheduling, internal search, meeting summarization.

Decided by: a standing rule plus one named reviewer's sign-off.

EXPEDITED · 1–2 WEEKS

Patient-adjacent. Advisory or influential impact, human in the loop, complete evidence pack, precedent already in the portfolio.

Decided by: a two-person subcommittee, escalating on disagreement.

FULL COMMITTEE · ONE CYCLE

Determinative impact, autonomous or on-the-loop control, novel modality, incomplete evidence — or anything touching consent.

Decided by: the full committee, as today.

The exact boundaries matter less than having them written down before the next submission arrives.

Four questions decide the lane

1 **Is the impact determinative?**

Does the output **become** the decision, or does a human decide with it? Determinative goes to the full committee regardless of anything else.

2 **Is there a human in the loop on every output?**

Not “can a human intervene.” Does one act on each output. On-the-loop and autonomous escalate.

3 **Does it touch consent, recording, or data leaving the boundary?**

If yes, full committee. This is where the current litigation is.

4 **Is the evidence pack complete?**

Incomplete submissions don't get a lane. They get returned. This is a routing rule, not a judgment.

Answer no to 1–3 and yes to 4, and it almost certainly does not need twelve senior people in a room.

THE HARD PART

The hard part isn't building the tiering. It's agreeing that some things don't get discussed.

Every committee has a member who believes their review is what makes the process safe. Sometimes they're right. Usually what makes it safe is that the evidence was complete and someone is watching after go-live.

The argument that works: **the committee's attention is the scarcest safety resource in the organization, and spending it on a scheduling optimizer is the thing that's actually unsafe** — because it's the reason the autonomous clinical model has been sitting in the queue for four months.

ONE THING THAT MAKES IT POLITICALLY SURVIVABLE

Publish what got rule-routed. Every quarter, a one-page list of everything that skipped the committee and why. Transparency about what you didn't discuss is what buys you permission to not discuss it.

Four numbers worth tracking

1 **Queue depth**

How many submissions are waiting right now. Most committees genuinely don't know this number, and it is the first one to get.

2 **Time in stage**

Intake → evidence complete → reviewed → decided. Almost always, one stage holds 70% of the elapsed time. Find it before you change anything else.

3 **First-pass completeness**

What share of submissions arrive complete against the published list. The best leading indicator you have of next quarter's throughput.

4 **Decisions per quarter, by lane**

If the rule-routed lane isn't carrying the majority of volume within two quarters, the gates are set too tight.

None of these require a system to start. A spreadsheet and one honest month will tell you where the constraint is.

A backlog is a design failure, not a governance failure.

Committees don't fall behind because the people on them have bad judgment. They fall behind because every submission gets the same path, and the path was designed when there were four tools a year.

Troy Bannister — Founder & CEO, Onboard AI

If your committee has solved this some other way, I'd like to hear it — I've only seen it work one way and I'd rather be wrong.