

A slow report can have three different causes.

Here is a fictional version of the reasoning we use in a workflow call. The useful output is a next check that separates the possible causes. Read the conversation and use the check yourself.

The question a founder brings

"Our assessment reports take too long. Should we buy an AI drafting tool?"

What they describe on the call: drafts wait for a senior reviewer, and some come back with similar corrections each time. That is their account of the situation. Nothing has been inspected yet, and the call does not pretend otherwise.

The conversation, worked

Delivery lead: "The reports take too long. The team has already tried AI for drafting."

Build n Bloom: "When you say too long, do people spend hours preparing them, or do finished drafts sit waiting for review?"

Delivery lead: "Both, I think. The partner sends them back."

Build n Bloom: "What comes back repeatedly: missing evidence, a conclusion the source does not support, or wording?"

Delivery lead: "Usually questions about where a statement came from."

Build n Bloom: "That gives us something to check; it does not establish the cause yet. Take one permitted recent report and note which returns asked for evidence already available to the preparer. Separately record how long an acceptable draft waited for review. If the returns are mostly missing references, try the traceable brief on the next comparable piece of work before buying another drafting tool."

Delivery lead: "And if that fixes it?"

Build n Bloom: "Keep the simpler change. If the repeated preparation still consumes material time, we can examine whether a separately scoped workflow is worth designing."

Three explanations that all fit the same symptoms

1. **Preparation** may be omitting the evidence the reviewer needs, so review turns into rework.
2. **Review** may be waiting on scarce senior time or unclear ownership, so acceptable work sits in a queue.
3. **Drafting** itself may genuinely be consuming the hours.

A drafting tool may help with writing effort, and if configured for the job it may also assist evidence preparation. But faster writing alone does not establish that incomplete preparation or review waiting has improved, and more than one explanation can be true at once. The useful move is not picking a favourite; it is finding an observation that tells the explanations apart.

The next check we'd agree

Take one recent, permitted assessment and record active evidence preparation, drafting, review and correction effort separately; record elapsed waiting for review alongside them, not as staff effort. Then read the stated reasons on returned drafts. This happens inside your firm, with your records; nothing is uploaded to us. One case will not prove where the bottleneck is, but it tells you what to look at before spending money.

How different answers change the move

- **Repeated missing-evidence returns:** examine the preparation-for-review standard before buying any tool. The free pack's traceable brief and evidence request are a no-cost way to test this.
- **Long waits on otherwise acceptable work:** examine review scheduling and ownership before deciding whether process change or tooling would help.
- **Genuinely heavy drafting effort:** then test drafting support properly, on held-out examples, comparing against reviewer-approved answers the tool cannot access.

Each answer points somewhere different, which is why buying a tool before running the check risks solving a problem you don't have.

A perfectly good ending

You leave with a specific observation to run and no reason to buy anything yet, from us or from a tool vendor. That is a successful call, not a failed one. If the check later shows repeated work that is worth designing around, the paid System Blueprint exists for that decision; it is a separate choice you would make with your own numbers in hand.
