

JUNE 5, 2026

The four pre-draft artifacts that keep AI accurate

Before AI drafts anything high-stakes, make it produce four things first. None of them is the draft – each one surfaces the model's reasoning while mistakes are still a one-line fix.

By **Dave Taylor**



Before you let an AI draft a high-stakes document, make it produce four things first: a source inventory, a conflict log, a missing-context list, and a duplicates report. None of them is the draft. Each one makes the model's reasoning visible so you can correct it while mistakes are still cheap – instead of discovering them inside a finished artifact that already reads as authoritative.

It's the same idea behind [structural hallucinations](#) and the [AI data room](#): AI accuracy is a data problem, not a prompt problem. These four artifacts are how you make the data problem visible before it becomes a credibility problem. The framing builds on the "AI data room" concept from [Nate B Jones](#). Here they are, in order.

1. The source inventory

A plain table listing every file in the workspace: its path, type, date, who produced it, and how far it can be trusted. This is the most boring artifact you'll ever ask an agent for, and the most valuable. It makes the model's judgement legible — you can see which documents it's treating as authoritative and which it's discounting, before any of that gets written into a draft. If the inventory looks wrong, you fix the data, not the document.

2. The conflict log

When an agent reads across a real set of files, it finds disagreements — the old plan says one number, the current one says another; two documents name different owners for the same task. A weak workflow lets the model quietly smooth those over, and the output reads beautifully while being quietly untrustworthy. A conflict log forces the disagreements into the open, so a human decides which source wins.

3. The missing-context list

Ask the agent to state plainly what it does *not* have. This matters more than it sounds, because missing material is the single biggest hallucination trap: if the model lacks a fact it needs, it will invent its way around the gap rather than stop. Made legible up front, those gaps become a short list of things to go find — instead of fabrications you have to catch later.

4. The duplicates report

Real folders accumulate near-copies: three versions of the same deck, two exports of the same spreadsheet. Left unflagged, an agent may silently blend or average them, corrupting its own reasoning with data it counted twice. A duplicates report makes it identify and quarantine the suspects so you can pick the real one.

Make the model's view legible before it drafts

The pattern across all four is the same: surface the model's view of the material *before* it commits to a draft. That's the whole move. Each artifact is a checkpoint where a human can intervene cheaply, while the cost of a wrong assumption is still a one-line fix rather than a rewrite.

You don't need new tools to start. On your next serious AI task, require these four outputs before any drafting begins. The first time you run it, the conflict log alone will usually surface something you didn't know was wrong in your own files — which tells you the

accuracy problem was never really about the AI.

