

The Hidden Cost of Doing It by Hand

Why the most valuable part of sales-and-use tax work is the part a spreadsheet throws away

When sales-and-use tax exposure work gets done in a spreadsheet or an internal tool and there is no software line item, the cost looks like zero. It rarely is — the real expense was never the software, but the expert judgment your team brings to each engagement and then leaves behind the moment it closes.

The value is the judgment, not the arithmetic

Calculating exposure, accruing interest and penalties, filling cells — that is the cheap part. The value is in the calls your people make: whether economic nexus is tripped and when, whether a product or SKU is taxable in a given state, when to recommend a VDA versus registering and back filing, how aggressive to be on a taxability position or a look-back period. That judgment is what a client pays for, and what separates one firm’s SUT practice from another’s. In a spreadsheet it is made once, lives in the cells and the expert’s head, and accrues nowhere — re-derived from scratch on the next engagement and gone the day a senior practitioner walks out the door.

What captures and compounds the expertise

An AI-native platform does not just store the work — it turns each call your experts make into something the whole practice can reuse. Four things make that possible, and a spreadsheet does none of them:

It learns from every decision — every nexus call, taxability determination, and VDA recommendation your experts accept, reject, or refine becomes a signal that sharpens the next one. The system gets measurably better each time an expert touches it. A spreadsheet cell teaches nothing — the next analyst starts cold.

It finds the right prior work by meaning, not keywords — a new fact pattern surfaces the comparable engagement — the taxability call or the nexus determination your team made last year — even when the wording is nothing alike. A spreadsheet can match a value; it cannot find “the time we worked out whether that bundled software charge was taxable in Texas.”

It grounds every answer in your own work and the actual rules — drafts and determinations come back tied to your prior workpapers and the controlling state authority, with citations attached, so the reasoning is reused and defensible rather than re-created from memory.

It is specific to the entity and the jurisdiction — the question, the taxpayer’s footprint, and the relevant state rules come together in every answer, so a determination fits this company in this state — exactly where a generic tool or a public chatbot fails.

Together these turn your team’s SUT expertise from something held in heads and spreadsheets into an asset the whole practice draws on — one that *gets sharper* with every engagement instead of resetting. And because the same platform runs the exposure estimate across 100% of transactions with the methodology documented as it is built, the workpaper also defends differently in a VDA or an audit — worth having, but almost table stakes next to expertise that compounds.

	Doing it by hand	On an AI-native platform
The expert judgment	<i>Made once, then left in a cell</i>	Captured with its reasoning as it happens
Finding a prior call	<i>Keyword or value lookup only</i>	Surfaced by meaning — the comparable nexus or taxability call
Grounding a determination	<i>Re-typed from memory</i>	Tied to your workpapers and the state authority, cited
Getting better	<i>No feedback loop</i>	Every nexus and VDA call sharpens the next
At turnover	<i>Walks out with the person</i>	Stays in the practice as firm IP

The question worth asking is not “*what can’t we do today.*” It is “**what are we letting evaporate every cycle that we could be compounding instead.**”