

Switchboard

From spreadsheets to a **modern claims platform**

How firms running on an aging internal system and a wall of spreadsheets get to one platform, without breaking a live case and without a single judgment call leaving human hands.


Your operation has a **shape**

Every settlement you administer is different: different class, different court, different fund. And every one of them runs the same line: notice plan, class list intake and NCOA, the case website and the call center, claims review with deficiencies and cure letters, opt-outs and objections, QSF funding, pro rata allocation, distribution and reissues, the 1099 run, escheatment, cy pres, and the post-distribution accounting at the end.

That sequence is the shape of the business. It's why the same team can run a 300-member wage-and-hour case and a 20,000-member data breach case: **the stations don't change, only the volume does.**

It's also where the trouble lives. The clean case rides the line end to end. The exceptions **fall off the line** into spreadsheets, inboxes, and memory: the returned check, the deceased claimant, the settlement funded in three tranches. If that reads like your week, this playbook was written for you.

○ THE CASE LINE · NOTICE THROUGH ESCHEATMENT



Three of them, drawn beside the station they leave. A firm running twelve cases at once is carrying dozens, and none of them are on the line where the reporting can see them.

THE PROBLEM, CONCRETELY

Follow one returned check


Here is the same piece of work in both worlds.

A check comes back: bad address. Today, that usually means: the envelope lands in the mail room, someone opens it, someone remembers which case it belongs to, someone opens the tracking spreadsheet, someone requests the address trace, and someone remembers to follow up. Six "someones," usually two people, across three weeks. And months later, the accounting the court requires has to explain what happened to that check. **So someone reconstructs the story** from the spreadsheet, the inbox, and memory.


On a platform, the same check: the returned status lands in an exceptions queue the day it bounces. The address trace is requested from the queue. A person approves the reissue (**that judgment stays human**), and the check goes back out with its whole history attached. When the court asks, the report already knows.

○ ONE CHECK · TWO WORLDS

TODAY · THROUGH PEOPLE'S MEMORY



ON A PLATFORM · THROUGH A QUEUE



Multiply this by every returned check, every deficient claim, every stale payment past 30 days, every deceased claimant, on every case you're running at once. **That's the gap between the two worlds.**

The field, **measured**

Most of what gets published about class actions is written for the lawyers filing them, not for the firms administering them afterwards. Read it from the operations side instead and five numbers explain why throughput is the problem worth solving:

55%

Of checks get cashed when payment goes out automatically; the rest become reissues and escheatment work

39%

Of settlements in the one district that requires a post-distribution accounting actually filed it

94%

Of the dollars in the 100 largest U.S. settlements ever went through just five firms

\$79_B

Moved through class and government settlements in 2025, a record, and \$238B across the four years to 2025

68%

Of class certification motions courts granted in 2025, up from 63% the year before; every class granted becomes a notice program and a claims period

Sources: FTC, "Consumers and Class Actions" (149-settlement study) · Prince, Walker & Dodson, N.D. Cal. settlement-distribution study · ISS Securities Class Action Services, Top 100 U.S. Class Action Settlements (administrator table) · Duane Morris Class Action Review 2026.

Read each one as a line in your own week. Uncashed checks are your reissue queue and your escheatment tracker. The missing accountings are the report your team assembles by hand. The concentration number means the firms with platforms are absorbing the volume. And the certification rate is the front of your pipeline: two motions in three that reach a ruling produce a class, and **every class produces a notice program, a claims period and a distribution** for somebody to administer.

Then there's the number that changed the job itself: suspect claims went from roughly **400,000 in 2021 to over 80 million in 2023**, single settlements now find half or more of their claims fake, and AI-generated claim fraud is a named problem in the field. More review per case, with no more fee to pay for it. **More cases, more claims, more fraud screening, against the same headcount.**

One structural fact sits under all five: **no off-the-shelf system exists for this work.** There is no "Salesforce for settlement administration." Every firm that wanted software had to build its own, which is why the field runs on internal apps of a certain age, with spreadsheets covering everything the app never learned to do.

If that describes your stack, it isn't a failure of your firm. **It's the condition of the whole niche.** What's changed is on the building side: custom software no longer requires an engineering department, which puts the platform the giants built for themselves within reach of a 30-person firm.

● THE WORRY


"Custom platforms are for the giants. We're not big enough for our own software."

● WHAT ACTUALLY HAPPENS

The giants built theirs when building took an engineering department. That's what's changed. A platform scoped to the **seven treasury families and the claims lane a firm your size actually runs** is a project measured in weeks. Mapping it is the whole first step, and that is the next chapter.

Map the work before you build **anything**

You wouldn't let a contractor knock down a wall in a house you live in without **drawing the house as it actually is**, including the wall everyone knows not to lean on. In an operation full of live cases and court deadlines, the map is that drawing.

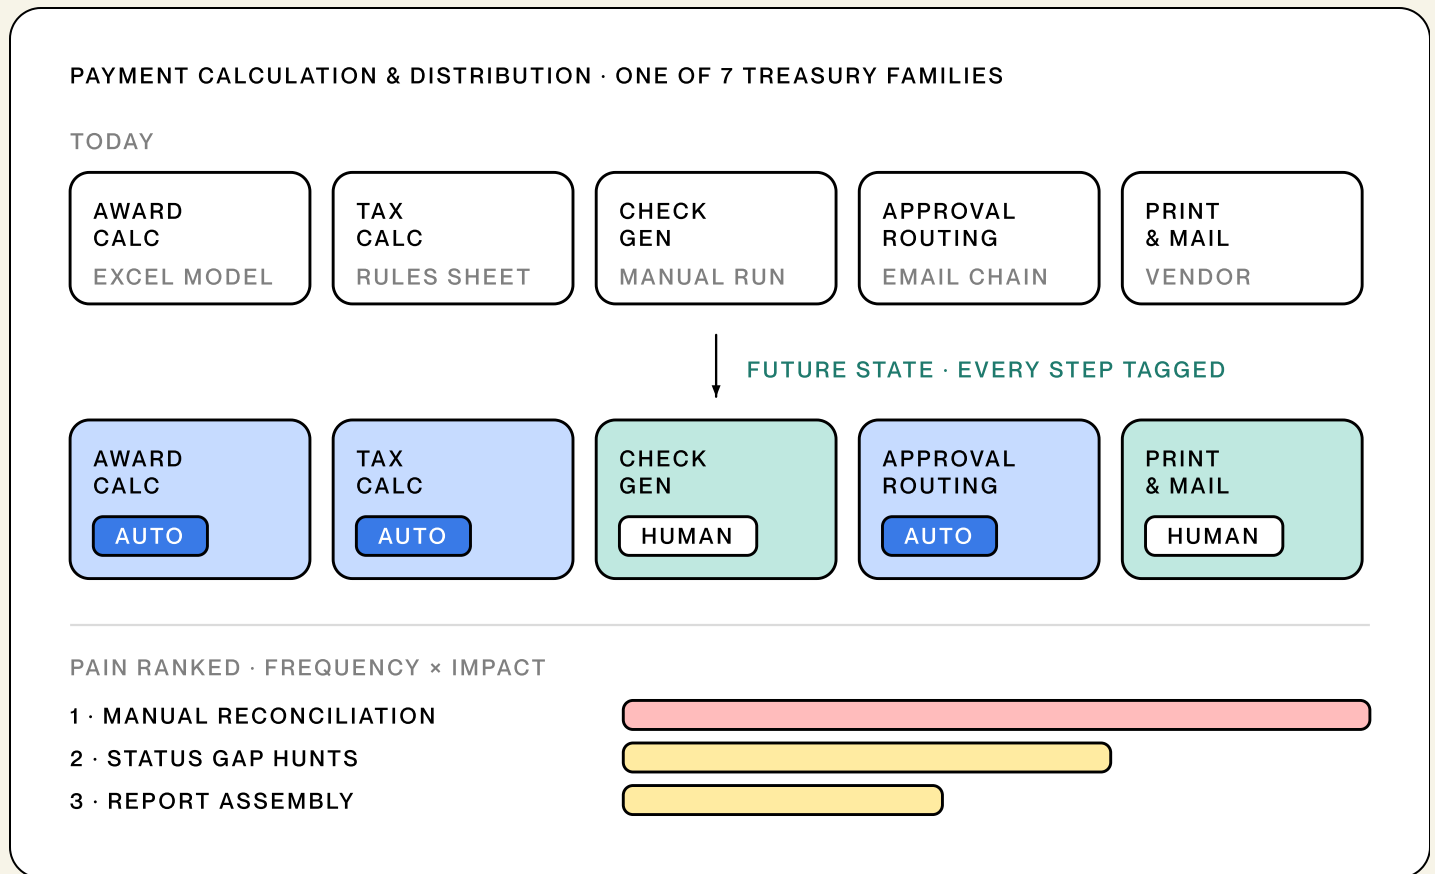


Mapping looks like this: interviews with the people who do the work, from the partner who prices bids to the person working the claims queue, and time spent watching the work happen. Every workflow gets drawn twice: **how it runs today, and how it should run**. Every future-state step gets a tag: automate, or keep human. The pains get ranked by how often they bite and how hard.

Two things surface on every map we draw. First, the work the org chart doesn't show: the tracking spreadsheet keeping a whole workflow afloat, the pricing model that lives in one person's head, the follow-up thread running on whoever remembers it. **Those get mapped too, and the build absorbs them**, instead of pretending they don't exist. Second, steps that turn up under more than one role: two teams describing the same part of a workflow, and no shared answer about which of them owns it. One operation we mapped had an entire workflow whose only job was finding records the other workflows had lost.


Below is **the map itself**, two pages of what you'd be handed in week two.

ARTIFACT 01 · A WORKFLOW MAP EXCERPT



One family from a real map, sanitized. Across the full treasury operation this covers 7 families and 26 steps; on the maps we've drawn for settlement operations, roughly 20 of the 26 take the blue tag.

● ARTIFACT 02 · THE DATA, ON ONE PAGE



When people say "the data model," they mean this: the eight kinds of records your business keeps, and how they connect. It comes out of the workshops we run with your team, and it becomes the platform's navigation.

● THE WORRY

"Our data is a mess. We'd have to clean everything up before we could even start."

● WHAT ACTUALLY HAPPENS

Messy data is a finding, not a blocker. The map turns "a mess" into a list: which records live where, which disagree, and which workflow keeps losing them. Cleaning up before mapping is how clean-ups fail: you'd be organizing the closet before drawing the house.

What automates, and what **never will**

Think **power steering, not self-driving**. You still hold the wheel; the strain is what disappears.

Every operator asks this first: **will AI be deciding claims, or paying people?** No. But in this business a bare "no" is worth very little, because the person asking has heard it from vendors who then shipped something else. So here is the map instead of the reassurance.

The line between machine work and human work is an output of the map you just saw: every future-state step carries its tag before anything gets built, and the tags follow one rule. **Wherever a physical act, a regulator, or a judgment call lives, a person stays.**

● THE SYSTEM DOES

- ◆ **Status tracking** on every payment, claim, and case, so nothing waits on someone remembering it
- ◆ **Claims intake and screening:** every filing scored for eligibility, duplicates and fraud, with the reason attached
- ◆ **The class member record:** name changes, address changes, deceased members and eligibility status, carried across a multi-year case
- ◆ **Reconciliation** against bank and vendor files, exceptions flagged the day they appear
- ◆ **Payment calculation** tier by tier, against the court-approved allocation
- ◆ **Self-service and triage:** the case website answers status, address and claim questions; what still arrives by email is routed with the case history attached
- ◆ **Report assembly:** the post-distribution accounting built from the record, not from spreadsheets
- ◆ **Tax prep:** information returns generated per class member, filing calendars that run themselves
- ◆ **Escheatment tracking** by claimant state, from 180 days outstanding, through to cy pres
- ◆ **Cure letter drafts** for deficient claims, queued for a person's review

● YOUR PEOPLE KEEP

- ◆ **Payment execution:** money moves only when a person moves it, in your bank portals
- ◆ **Check generation** and final print & mail
- ◆ **Eligibility judgment** on reissues, deceased claimants, and edge cases
- ◆ **Every claim the screen flags:** the model sorts and explains, a person clears or denies
- ◆ **Pricing:** bids stay a human call, now with the bid history in front of them
- ◆ **Everything court-facing:** declarations, approvals, filings
- ◆ **Government submissions:** state escheatment portals, IRS filing, remittance

Reconciliation

Checking that what you sent matches what the bank says happened. Today that means CSV files and eyeballs; on a platform it runs continuously.


Exceptions queue

One list that catches every returned check, missing tax ID, and stale payment the day it happens, instead of an inbox and a memory.

Audit trail

Who did what, when, and why, written down automatically on every record. The answer to any question the court asks later.

○ ONE PAYMENT · THREE HUMAN GATES



Every gate waits for a person. The system never moves money. It tracks, reconciles, assembles, and drafts; people decide and approve, and every action lands on the audit trail: actor, time, reason.

● THE WORRY

"Automation here means a chatbot talking to claimants, and AI making calls it shouldn't."

● WHAT ACTUALLY HAPPENS

Almost none of what pays off in this business is a chatbot. It's **queues, statuses, reconciliation, and reports**: the unglamorous software your day runs on. AI earns a place inside that (drafting cure letters for review, catching duplicate claims), never in front of it.

Where the margin **hides**

Your fee comes out of the fund, the court sees every dollar of it, and you won it by competitive bid. You can't raise the price. So every hour of manual reconciliation, claims chasing, and report assembly is paid out of a fee that was fixed the day you won the case. **Throughput per person isn't an efficiency metric in this business. It is the margin.**

And those costs aren't private. Under the procedural guidance courts now follow, they get filed line by line: administrative cost, cost per claimant, even how the administrator was selected.

It cuts the other way too: the firm that can promise faster reporting and cleaner accountings, and show the numbers to class counsel, **wins more bids and can afford to price sharper.**

The arithmetic, worked on a mid-size operation. Swap in your own numbers; the interactive version of this playbook and its workbook run them live.

Worked on a 20-person case team, the calculator's conservative defaults:

20 people × 150 hours a month × 25% absorbed by the platform = **750 hours back per month**

750 hours ÷ 150 hours per person, with 20% annual case-volume growth = **6 roles you don't have to add**

6 roles × \$60K loaded cost = **\$360K a year in avoided hiring**, which pays a \$250K build back in **about 8 months**

Both costs are placeholders to swap for your own. The mapping phase replaces all of this with the hours your team actually logs, workflow by workflow.

Find yourself on the **curve**

Every firm that runs settlements sits on this four-stage curve, and none of the stages is a judgment: as chapter 2 showed, the whole niche climbed here together. **Find the Monday that sounds like yours**; each stage carries its next move, and it is never "buy everything at once."

● STAGE 1

Everything is a spreadsheet

MONDAY LOOKS LIKE

The allocation model, the check log, the escheatment tracker: all spreadsheets, each owned by whoever built it. The case history lives in inboxes.

THE NEXT MOVE

Don't buy software. Map the one workflow that hurts most; the map is what tells you what to build or buy.

YOU'VE OUTGROWN IT WHEN

A case can run without asking the person who built the spreadsheet.

● STAGE 2

A system that covers half

MOST OF THE FIELD IS HERE

MONDAY LOOKS LIKE

The internal system handles the clean case. Everything the system never learned (the exceptions, the reconciliation, the report) lives in the spreadsheets around it.

THE NEXT MOVE

Map the halo of spreadsheets around the system. That's where the hours go, and the map prices exactly what replacing them is worth.

YOU'VE OUTGROWN IT WHEN

You can answer "who has been paid?" without opening a spreadsheet.

● STAGE 3

One platform

MONDAY LOOKS LIKE

Exceptions in queues, statuses tracked end to end, the court report assembled from the record. The judgment calls are still yours. That's the design.

THE NEXT MOVE


Deepen the automation, workflow by workflow. Your team's hours start going to judgment instead of tracking.

YOU'VE OUTGROWN IT WHEN

The court report is a button, and month-end closes without a fire drill.

● STAGE 4

AI on the rails



MONDAY LOOKS LIKE

Cure letters drafted for review, duplicates caught at intake, anomalies flagged before they age. Every AI feature rides the platform's clean rails, and every action is audited.

THE NEXT MOVE

This stage is only reachable from Stage 3. AI on top of spreadsheets is noise. AI on top of rails compounds.

YOU'RE DOING IT RIGHT WHEN

Every AI draft can show the record it came from.

You climb this curve **one mapped workflow at a time**, never by buying everything at once, and the last chapter is what that first workflow looks like.


Your first two weeks **with us**

The renovation rule applies: **nobody tears out the kitchen on day one**. Your current system keeps running until the new one has earned trust on live work.

Starting with us is smaller than most operators expect. What we ask for: **read access to how you work today, and workshops with the people who do the work**. No migration, no downtime, nothing touched in production.

From there, the sequence this playbook has walked through is the one we run: interviews and observation in the first two weeks, then the map: every workflow drawn current and future state, every step tagged auto or human, the pains ranked. And in the weeks after mapping wraps, **a clickable prototype of your platform**: your statuses, your vocabulary, a believable case built on sample records rather than your live data. You see the movie before you make the movie: software you can click, not a slide deck about software.

○ FROM READ ACCESS TO WORKING SOFTWARE



The person marks the steps your team carries (the workshop weeks); the spark marks ours. **Your current system runs untouched the whole time.**

Five things come out of those two weeks, and all five are yours to keep whether or not you build with us.



The map

How work actually moves today, and where it breaks, on paper and in plain English



The opportunity map

What's worth building, ranked by what it returns



The tech audit

Your stack, the integrations, and what production requires



The prototype

Working software on your workflows: click it, react to it, redirect it



The build plan

Phased and sequenced from your reactions, specific enough to build from

Everything above is **yours**. The map, the plan, the prototype: built around your operation, owned by your firm.

● THE WORRY

"This is a rip-and-replace with a year of migration, and our live cases can't survive that."

● WHAT ACTUALLY HAPPENS

Nothing migrates until it's proven. The first build is scoped to the workflow that hurts most, in weeks, not years, and it runs alongside your current system, which keeps handling live cases until the new platform has earned each workflow, one at a time.

If this looks like your operation, **let's talk**

Thirty minutes. No pitch, no deck. A conversation about where your firm sits on the curve, and what mapping it would look like. **Book a call: withswitchboard.com/book-a-call**

This playbook also lives as an interactive page, with a calculator for your own numbers, a stage-marking that routes the reading, and a take-home workbook for your team: **withswitchboard.com/playbook/settlement-administration**