Nexusflow INC.

RESEARCH REPORT | 2026

Stale by Thursday

The hidden cost of manual spreadsheet work, and how to get your week back

nxflo.com | Vancouver, Canada

November 2026

Contents

Executive summary . . . 3
Anatomy of a report that is stale by Thursday . . . 4
How much time is it, really? . . . 5
Which documents cause the most pain . . . 6
Data teams feel it most, but the numbers disagree . . . 7
What it costs . . . 8
The accuracy problem . . . 9
Why small errors become big ones . . . 10
The human cost . . . 11
Why people do not automate . . . 12
What people say online . . . 13
Choosing what to automate first . . . 14
Finance and accounting . . . 15
Operations, supply chain and logistics . . . 16
Sales, marketing and customer support . . . 17
HR, administration and operations support . . . 18
Facilities, property and industrial operations . . . 19
The idea behind the fix: ETL in plain English . . . 20
The automation ladder . . . 21
Scoring your candidates . . . 22
The payback maths . . . 23
A 30-60-90 day plan . . . 24
Keeping automation safe . . . 25
Build, buy or hire? . . . 26
Nexusflow free tools: a concept . . . 27
How Nexusflow can help . . . 28
Glossary and common questions . . . 29
Sources and notes . . . 30

How to read this report. Each page stands on its own. Part 1 and 2 describe the problem; Part 3 shows where it appears by function; Part 4 is a practical guide to the fix; Part 5 covers what Nexusflow offers. Several surveys were vendor-commissioned and online posts are anecdotes, so read figures as signals of direction, not guarantees. Our own illustrations are labelled.

IN ONE PAGE

Executive summary

Every week, a very large number of people download the same files, clean them, merge them and rebuild the same report by hand. By the time it is finished, the data has already moved on. This report brings together survey data, academic research and online discussion to show what that routine costs and what to do about it.

9+ hrs
per week on manual data entry, on average (Parseur, 2025, n=500)

94%
of 88 audited real-world spreadsheets contained errors (Panko review)

46%
of workers have never used any tool to automate data entry (Parseur)

Six findings

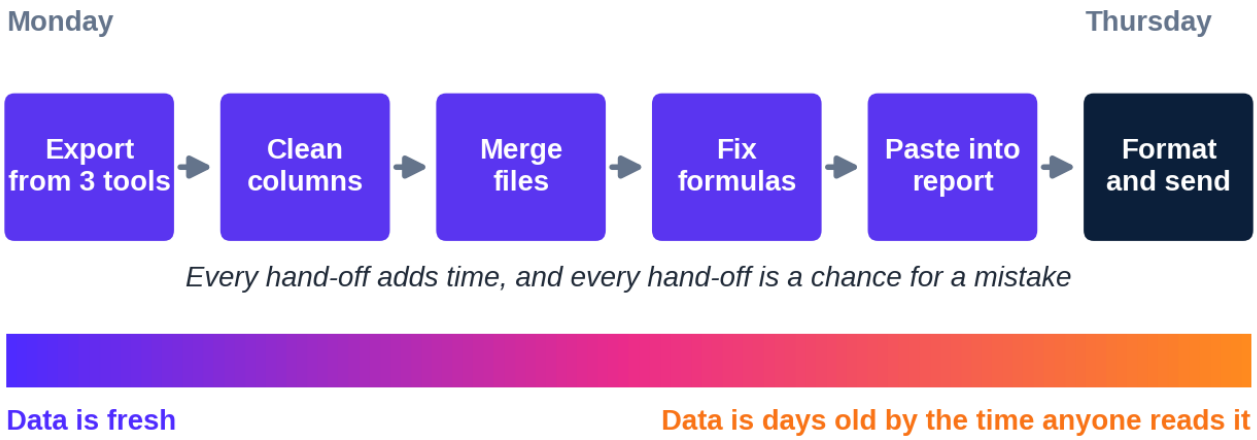
- **Time.** Workers report spending more than nine hours a week moving data between emails, PDFs, spreadsheets and scanned documents.
- **Money.** One survey puts the cost at $28,500 per employee per year; an IDC study for Alteryx estimated about $12,000 per user per year in duplicated effort.
- **Accuracy.** Audits of real spreadsheets find errors in the vast majority, and even carefully built ones keep errors in about 1% or more of formula cells.
- **Freshness.** Up to 8 hours a week can go to redoing work because the underlying data was refreshed, so reports are out of date on arrival.
- **People.** 56% report burnout from repetitive data tasks, and about half say it causes costly errors, delays or lost opportunities.
- **Opportunity.** Adoption is low, and among companies that do automate, 96.5% report a significant workload reduction.

What we recommend: pick the single weekly chore that annoys people most, time it for two weeks, automate that one thing, check the new output against the old for several cycles, and only then move to the next. Part 4 shows how.

THE PROBLEM

Anatomy of a report that is stale by Thursday

Most reporting still begins with an export. Someone downloads files from two or three systems, tidies the columns, merges them, fixes the formulas, pastes the result into a template and sends it. Every step is small. Together they eat the week.



Illustrative workflow of a typical manual weekly report. Not survey data.

Three costs hide inside the routine

- **Time.** The work is repetitive and predictable, which is exactly why it is invisible. Nobody budgets for it, so nobody notices it growing.
- **Errors.** Each copy, paste and manual formula edit is a chance for a mistake that nobody sees until a number looks wrong in a meeting.
- **Freshness.** The slower the rebuild, the older the data. A report finished on Thursday describes Monday, and decisions are made on yesterday's picture.

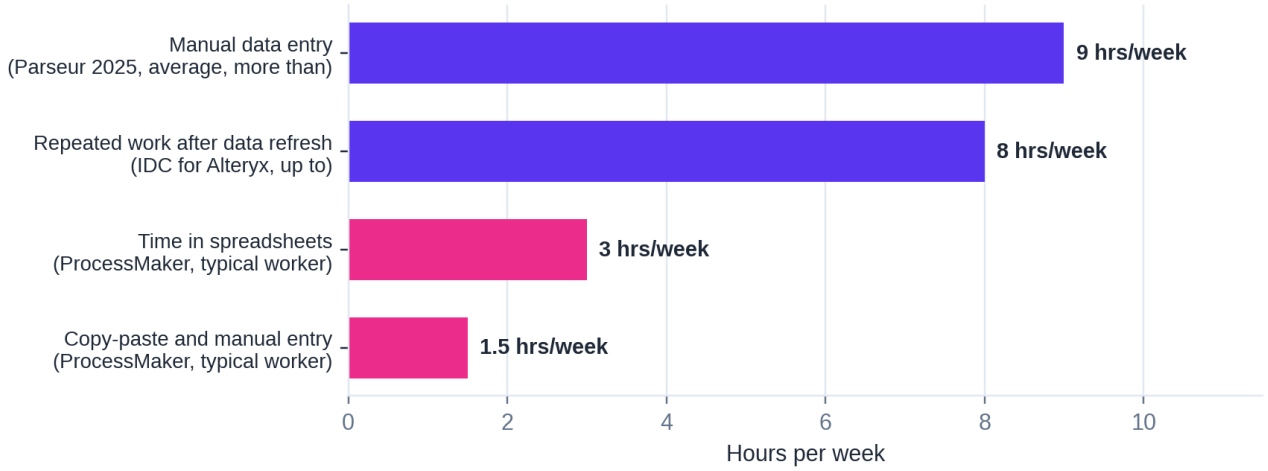
> The core idea of this report: the problem is not that people are bad at spreadsheets. It is that a spreadsheet is a poor machine for doing the same job over and over, and a better machine exists.

The rest of Part 1 puts numbers on each of the three costs: how much time, how much money, how many errors, and what it does to the people involved.

THE PROBLEM

How much time is it, really?

Depends who you ask and what you count. Published figures range from about an hour and a half to more than nine hours a week, and they measure different things in different populations.



Sources: Parseur/QuestionPro 2025 (self-reported, 500 US professionals); IDC for Alteryx (500+ spreadsheet users, advanced users); ProcessMaker (observed activity from opt-in enterprise clients). Figures are not additive.

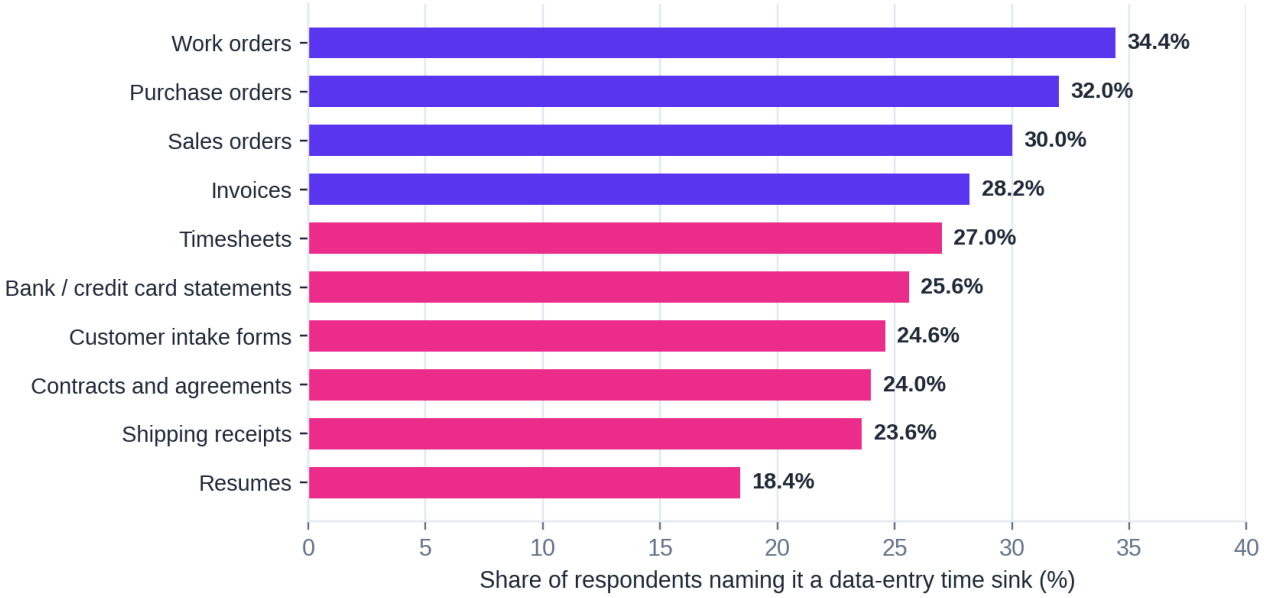
- **Parseur (2025)** asked professionals in finance, operations, IT, support and HR how long they spend transferring data from emails, PDFs, spreadsheets and scans into digital systems. The average was more than nine hours a week.
- **IDC for Alteryx** found that advanced spreadsheet users spend up to eight hours a week repeating work because the data was refreshed, out of about 26 hours a week in spreadsheets overall.
- **ProcessMaker** measured activity rather than asking. Across more than four million data points, the typical office worker spent about 10% of their time on manual data entry, around three hours a week in spreadsheets, and about an hour and a half copying, pasting or keying data into business applications. The average enterprise employee made over a thousand copy-paste actions a week.

How to use this: do not borrow a number. Time your own chore for two weeks. If it takes 90 minutes every Monday, that is 75 hours a year, whatever the surveys say.

THE PROBLEM

Which documents cause the most pain

The paperwork people say they hand-key most often is operational: orders, invoices, timesheets and statements. These are exactly the files that arrive as PDFs, emails and exports.



Source: Parseur/QuestionPro, July 2025, 500 US professionals. Respondents could name more than one document type, so shares do not sum to 100%.

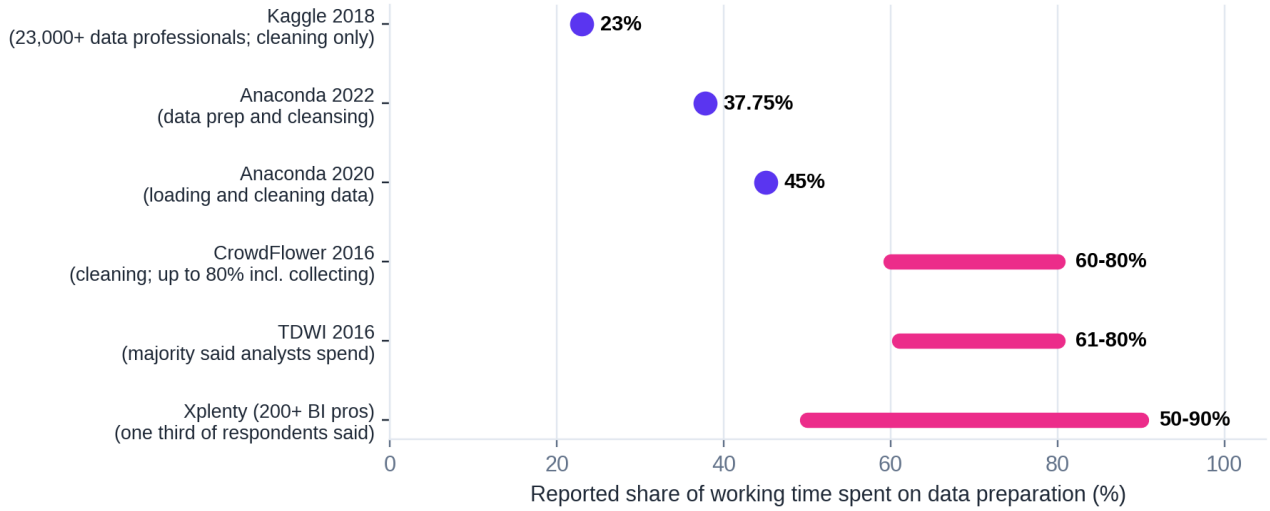
- **Work orders lead** at 34.4%, followed by purchase orders (32%), sales orders (30%) and invoices (28.2%).
- **Timesheets, bank and credit card statements and customer intake forms** each sit between roughly a quarter and a little over a quarter of respondents.
- **Specialised documents** also surfaced in the open answers, including educational grades, patient records, prescriptions and inventory reports.
- **The heaviest users** (20 or more hours a week) were disproportionately in IT and finance, who tend to sit in higher pay brackets, which raises the cost of the time lost.

> For a first free tool or a first client project, the safe bet is a document type from the top half of this chart. Real demand is already there.

THE PROBLEM

Data teams feel it most, but the numbers disagree

The claim that analysts spend 80% of their time cleaning data is repeated everywhere. The evidence is messier: surveys range from under a quarter to nine-tenths, depending on who is asked and how.



Sources: Kaggle 2018 survey; Anaconda State of Data Science 2020 and 2022; CrowdFlower 2016; TDWI 2016; Xplenty survey of 200+ BI professionals. Populations and definitions differ.

- **Trifacta (2018, n=294)** found 65% of respondents spend at least half their time preparing data, 92% would rather do other analytic work, and 58% believed the time spent costs more than it delivers. Excel was the most-used preparation tool (37% of analysts, 30% of IT staff).
- **DAA/TMMData** found 38.7% of data professionals spend more than 20 hours a week accessing, blending and preparing data.
- **Another survey of analysts** found 42% spend one to five hours a week cleaning and preparing data, 40% up to ten, and about one in twenty more than ten.
- **Xplenty** found 55% of BI professionals named integrating data from different platforms as their biggest challenge.

> Even the conservative end, roughly a quarter of the working week, is a large share of the time of well-paid people. The disagreement is about size, not direction.

THE PROBLEM

What it costs

Time is the visible cost. Multiplied by salary, it becomes the one finance teams notice.

$28,500
per employee per year in manual data-entry costs (Parseur, 2025)

$12,000
per user per year in duplicated effort (IDC for Alteryx)

$60B
estimated annual US productivity loss from repeated spreadsheet work (IDC)

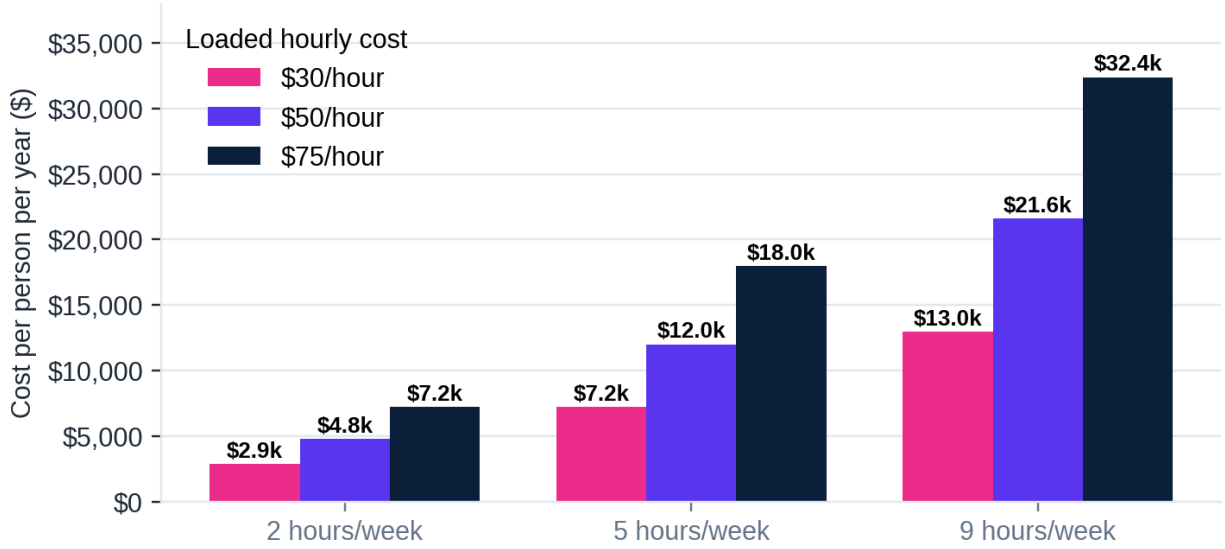


Illustration by Nexusflow: hours per week x hourly cost x 48 working weeks. Pick your own hours and loaded hourly rate (salary plus benefits and overhead).

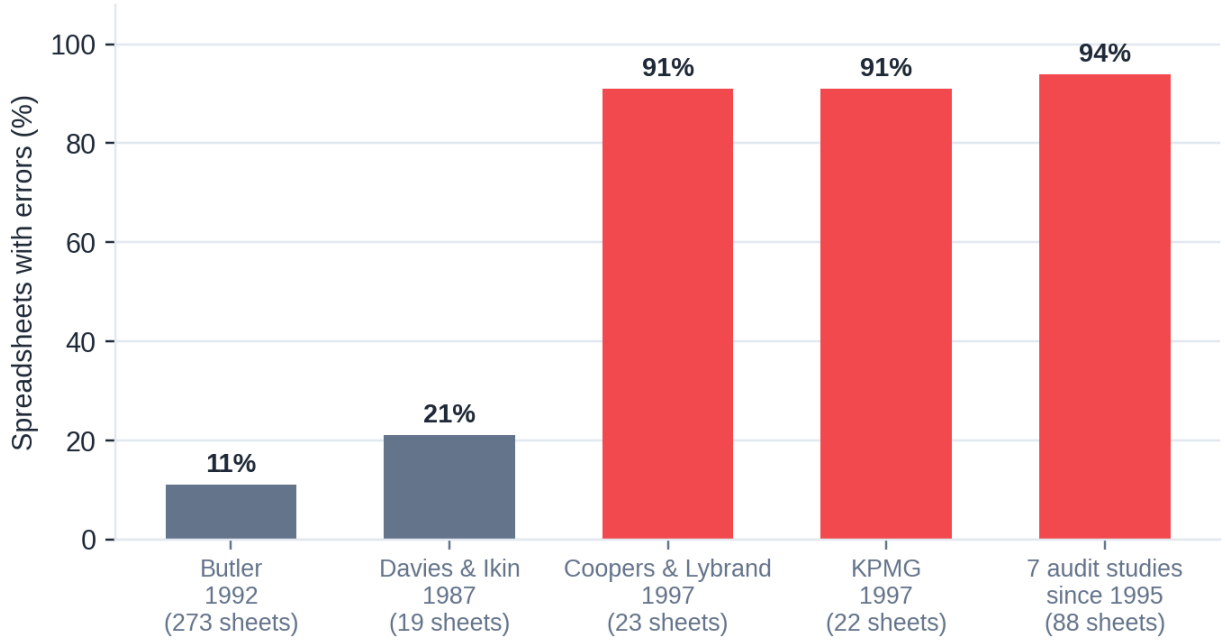
- Parseur notes that people who spend 20 or more hours a week on data entry are concentrated in IT and finance, with pay brackets of roughly $50 to $90 an hour.
- The cost estimates above are vendor-commissioned and use different assumptions, so they are best read as an order of magnitude.
- **Hidden costs not shown:** errors that have to be corrected later, decisions made on stale data, and the cost of people leaving because the work is tedious.

A quick test for your own team: hours a week, times an honest hourly cost, times 48. If the answer would pay for an automation project in a few months, the case is already made. Page 23 shows the payback maths.

THE PROBLEM

The accuracy problem

Manual spreadsheets do not just waste time. They are wrong more often than most people expect, and research on this goes back decades.



Sources: field audits compiled by Raymond Panko (University of Hawaii). Sample sizes are small and methods improved after 1995, when 94% of 88 spreadsheets audited across seven studies contained errors. Some studies counted only serious errors.

- In laboratory experiments covering about 1,170 spreadsheets built by 998 people, 51% contained errors even though most were only 25 to 50 cells in size.
- Panko notes that spreadsheets, even after careful development, contain errors in about one percent or more of formula cells, so large workbooks with thousands of formulas probably hold dozens.
- Experience did not help much: when MBA students with at least 250 hours of spreadsheet development were compared with novices, error rates were not significantly different.
- The Coopers & Lybrand audit only counted an error if it moved a bottom-line value by 5% or more, and still found one in 91% of sheets.
- A well-known real case: Toronto-based TransAlta took a $24 million charge after a cut-and-paste error in an Excel spreadsheet caused a bidding mistake.

> The aim of automation is not to make people faster at copy and paste. It is to remove the copy and paste.

THE PROBLEM

Why small errors become big ones

A one-percent error rate sounds harmless. Over a few hundred formulas it is almost a certainty that something is wrong.

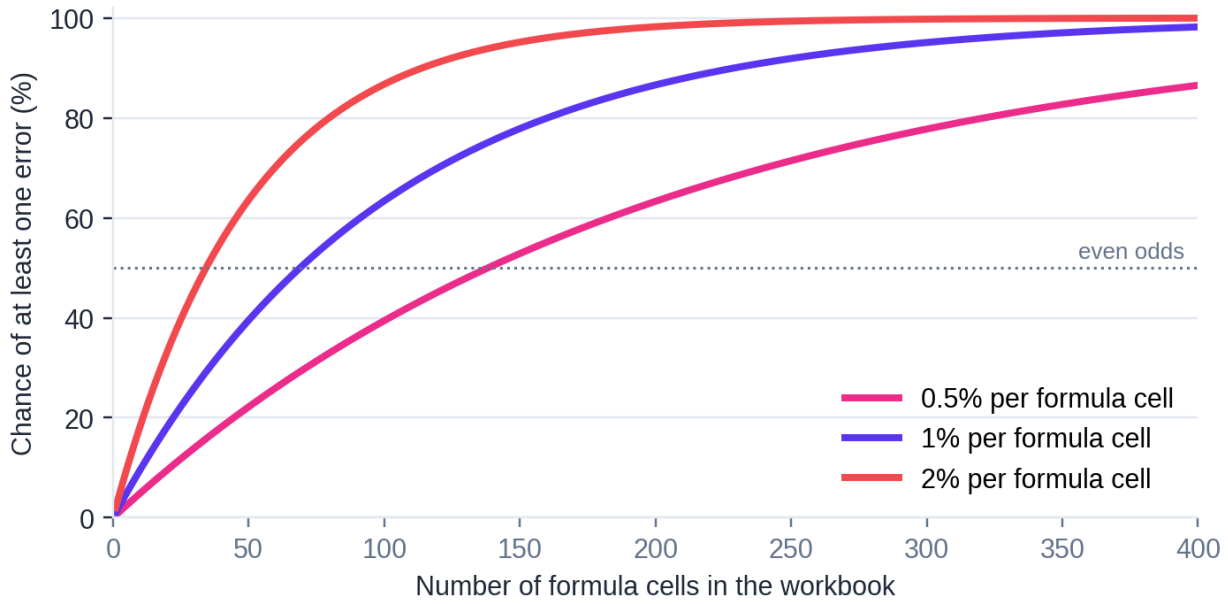


Illustration by Nexusflow: probability of at least one error = 1 - (1 - p)^n, assuming errors are independent. In real workbooks errors cluster and propagate, so this is a simplified model.

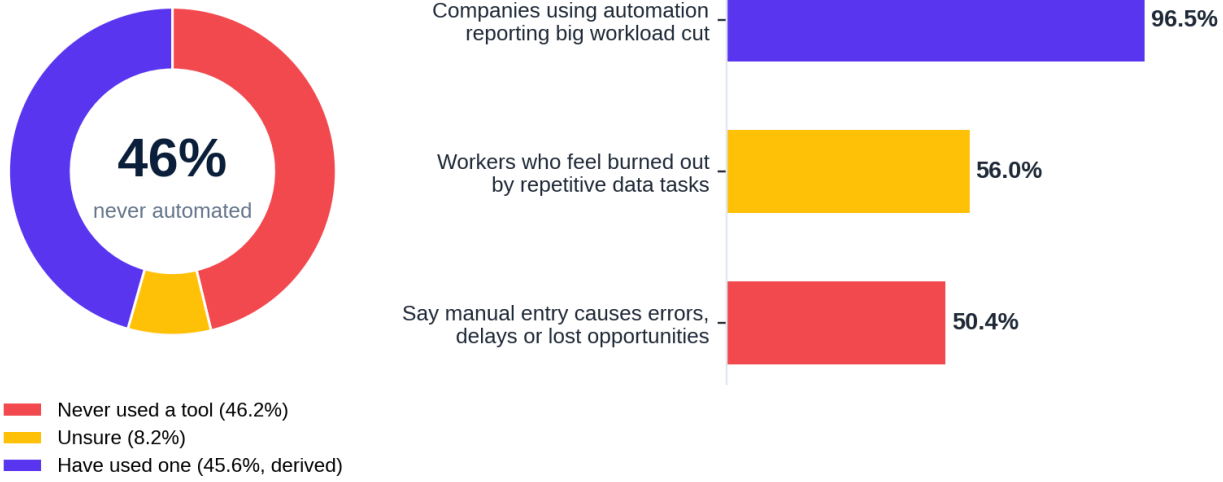
- At a **1% error rate**, a workbook with 100 formulas has about a 63% chance of containing at least one error. At 300 formulas it is about 95%.
- Spreadsheets chain calculations together, so one wrong input feeds everything that depends on it.
- Most errors are not dramatic. They are a wrong range, a stale link or a pasted value that should have been a formula, which is why they survive review.
- People are consistently overconfident about their own spreadsheets, which Panko found in every empirical study he reviewed.

Practical implication: every manual touch point you remove is a place an error can no longer enter. Automation does not make a bad process good, but it makes a good process repeatable.

THE PROBLEM

The human cost

The most overlooked cost of repetitive reporting is what it does to the people who do it.



Source: Parseur/QuestionPro, July 2025 (n=500). 'Have used one' is derived as the remainder after the never-used and unsure shares.

- **56%** of respondents said repetitive data tasks leave them burned out or frustrated.
- **50.4%** said manual data entry has led to costly errors, delays or lost opportunities.
- **46.2%** have never used a tool to automate data extraction or entry, which Parseur attributes to a lack of awareness and internal advocacy.
- **96.5%** of respondents whose companies do use automation reported a significant reduction in workload.
- Asked where they would redirect saved time, employees named strategic planning, improving the customer experience and revenue-focused work.

> Automation here is less about replacing people and more about giving them back the part of the job they were hired to do: analysis, judgement and action.

THE PROBLEM

Why people do not automate

If the case is this clear, why does the routine survive? Surveys and online discussion point to the same handful of reasons.

Barrier	What people say	How to address it
Accuracy	Will the automated version get the numbers right?	Run old and new side by side for several cycles and reconcile every total before switching.
Data security	I do not want to upload company data to a stranger's website.	Prefer tools that run locally or inside your own environment, and document where data goes.
Complexity	I am not technical and this sounds like a project.	Start with one report. Keep the first build small enough to explain in five minutes.
Not that bad	It only takes 20 to 30 minutes, I will get to it next week.	Multiply by 52. Twenty-five minutes a week is over 20 hours a year, for one chore.
Awareness	Nobody ever showed me it could be done.	Show a before-and-after on a real report from the team.
Payoff doubt	It might take longer to automate than to do by hand.	True for rare tasks. For weekly or urgent ones, the build usually pays back within months.

One participant in an r/excel discussion made the point well: even when an automation takes longer to build than a single manual run, it can still pay off if the task keeps recurring and has to be finished fast. Another noted that building automation produces knowledge as well as output, which makes the next one faster.

Parseur's survey adds a framing that matters for how you sell this internally: the three main concerns of hesitant workers are accuracy, data security and complexity. Answer those three before you talk about features.

VOICES

What people say online

Reddit and Product Hunt threads about automation keep returning to the same few chores. These are anecdotes, not statistics, but they show the language people use.

Theme	What comes up	Where
Weekly reporting	The most common answer to ‘what did you automate and never go back?’. One poster spent about two hours every Friday pulling client reports before automating it. Another described it as collecting the same items from the same places every week.	r/automation
Combining exports	Pulling data out of several tools and merging it into one spreadsheet is a job people keep promising themselves they will automate ‘next week’.	Product Hunt, r/automation
Trimming big reports	Teams download large reports from a BI tool every week and cut them down to what their group needs.	r/excel
Daily intake forms	Small daily forms that someone retypes into another system.	r/automation
Cleaning messy sheets	Sorting and restructuring untidy spreadsheets disappears hours without anyone noticing.	r/automation
Relief afterwards	People describe regaining mental space and losing the Monday dread once the chore was automated.	r/automation

What the pattern tells us

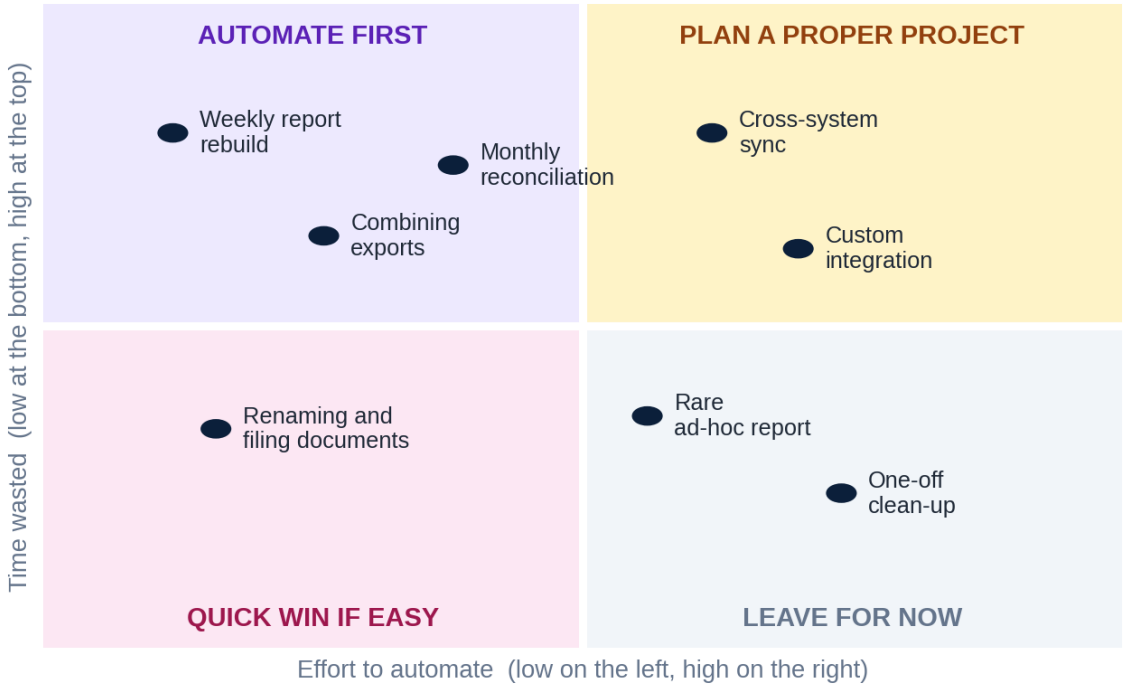
- The pain is rarely a hard technical problem. It is boring, regular and slightly too small to get budget, which is why it persists.
- People describe the task by its steps (‘export, clean, merge, paste’), not by the word ETL. Marketing should use their words.
- Many ask for help in public. That makes these communities a good place to listen before selling anything.

Method note: thread content was paraphrased from public posts in r/automation, r/excel and Product Hunt, accessed through archive mirrors. Individual posts may have been edited or removed since.

VOICES

Choosing what to automate first

Not every annoying task is worth automating. Two questions sort most candidates: how much time does it waste, and how hard is it to automate?



Illustrative framework by Nexusflow with example tasks placed for discussion. Positions are judgement calls, not measured data.

- **Automate first:** frequent, time-consuming, and built from predictable steps. The weekly report is the classic example.
- **Plan a proper project:** high waste but needs several systems to talk to each other. Worth doing, but scope it properly.
- **Quick win if easy:** small but trivial to fix. Do these when you have a spare hour.
- **Leave for now:** rare, or so awkward that the build would never pay back.

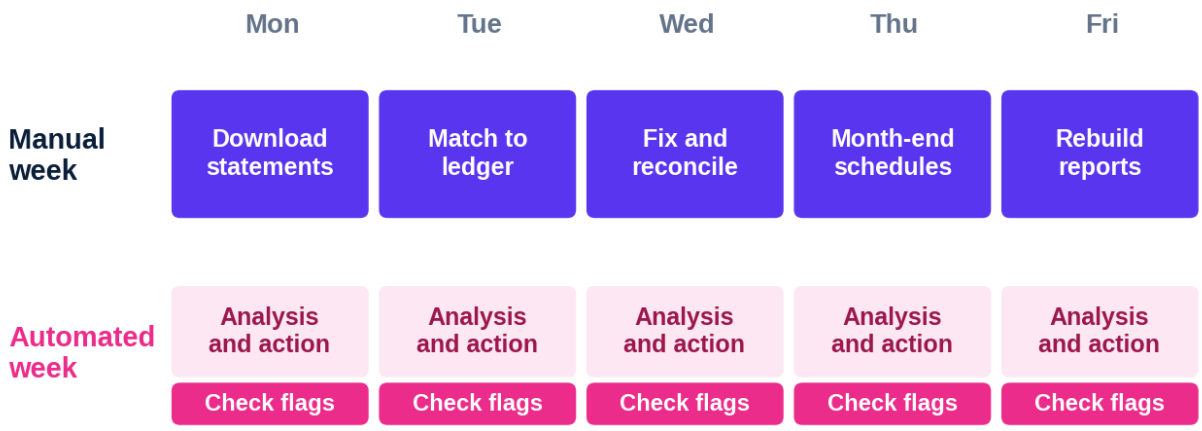
> Place your own five most annoying chores on this grid before you decide. The act of placing them usually makes the first choice obvious.

WHERE IT HURTS

Finance and accounting

Finance teams sit at the end of every other team's data, which makes them heavy users of exports and the first to feel stale numbers.

The weekly chore	Download bank and card statements, match them to ledger entries, rebuild the month-end schedules, and refresh management reports.
Why it hurts	Everything is a copy of something else. A single mistyped invoice or misaligned range ripples into reconciliations and reports, and month-end is a deadline that cannot slip.
What automated looks like	Statements and exports land in a folder, are read and matched automatically, exceptions are flagged for a human, and the standard reports refresh on a schedule.
A good first project	Bank or credit card statements to a clean, categorised table with a list of items that did not match.



Illustration, not measured data

Illustration by Nexusflow of how a typical week can change. Chores shown are examples.

Start in three steps

- Time one month-end cycle and list every step.
- Pick the statement format you receive most often.
- Automate matching for that one account and compare four cycles before switching.

Evidence: invoices (28.2%), timesheets (27%) and bank or credit card statements (25.6%) were all named by more than a quarter of respondents as manual data-entry burdens. IT and finance were over-represented among heavy data-entry workers.

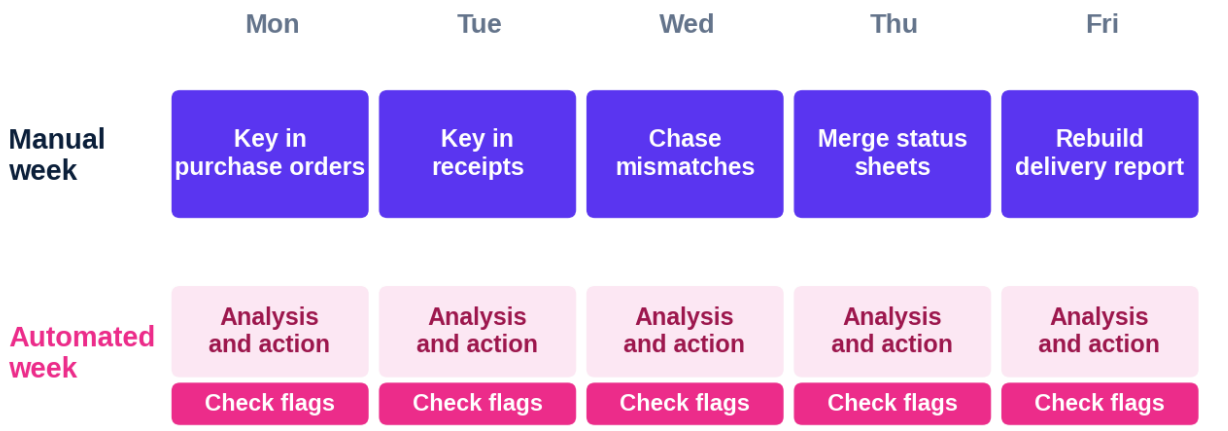
Source: Parseur/QuestionPro 2025. Task descriptions are Nexusflow's general illustration.

WHERE IT HURTS

Operations, supply chain and logistics

Operations runs on orders, receipts and status updates, many of which arrive as PDFs, emails or spreadsheets from outside partners.

The weekly chore	Key purchase orders, shipping receipts and work orders into internal systems, then rebuild status and delivery reports from several sources.
Why it hurts	Volume is high, formats vary by supplier and errors are costly: a wrong quantity or address becomes a reshipped order or a missed delivery.
What automated looks like	Incoming documents are read automatically into a standard layout, checked against rules, and pushed into the order or inventory system. A live status view replaces the weekly summary.
A good first project	Purchase orders or shipping receipts from the two or three busiest suppliers into a standard table.



Illustration, not measured data

Illustration by Nexusflow of how a typical week can change. Chores shown are examples.

Start in three steps

- List your two or three busiest suppliers and their document formats.
- Automate reading one document type into a standard table.
- Add a rule check, such as quantity or address mismatch, before anything reaches the system.

Evidence: work orders (34.4%), purchase orders (32%), sales orders (30%) and shipping receipts (23.6%) are all high on the list. One distribution-industry analysis cited by a vendor describes an order taking 8 to 10 minutes to key by hand versus 1 to 2 minutes of review after automation, as an example, not a universal figure.

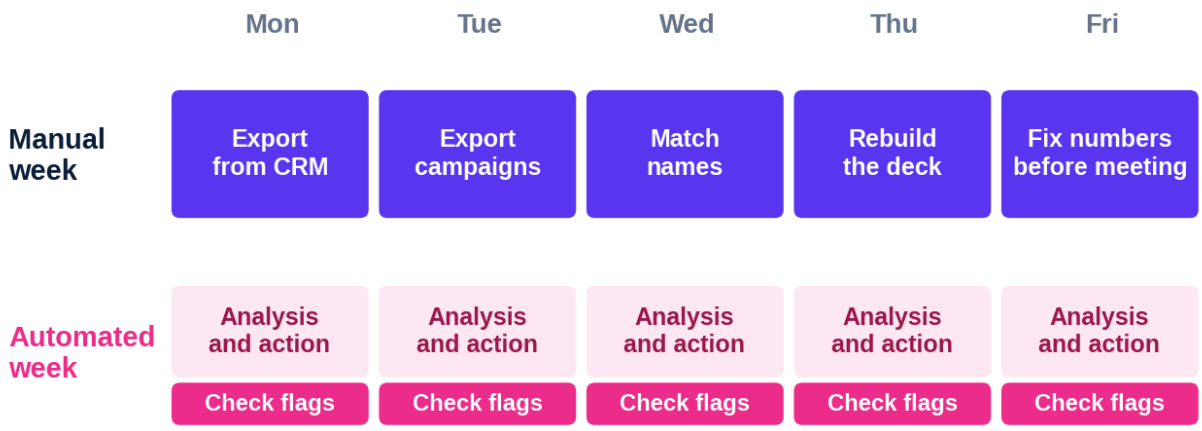
Sources: Parseur/QuestionPro 2025; a February 2026 industry article summarising the Parseur survey for distributors.

WHERE IT HURTS

Sales, marketing and customer support

Commercial teams report on activity: leads, campaigns, tickets and revenue. The data lives in different tools that do not agree.

The weekly chore	Export from the CRM, the ad platforms and the support desk, line the columns up, and rebuild the weekly performance deck.
Why it hurts	Every tool names things differently, so someone spends hours matching them. Client reporting in particular tends to land on Friday afternoons.
What automated looks like	A scheduled job collects each source, standardises the names, and refreshes one dashboard. The weekly meeting starts from the same numbers every time.
A good first project	A single weekly report that merges the CRM export with the campaign export, refreshed automatically.



Illustration, not measured data

Illustration by Nexusflow of how a typical week can change. Chores shown are examples.

Start in three steps

- Write down the five numbers the weekly meeting actually uses.
- Automate collecting those from each tool into one table.
- Refresh one shared dashboard on a schedule and retire the deck rebuild.

Evidence: sales orders (30%) appear high on the Parseur list. In online discussions, client reporting is repeatedly named as a weekly chore that people automated and never went back to, including one poster who spent about two hours every Friday on it.

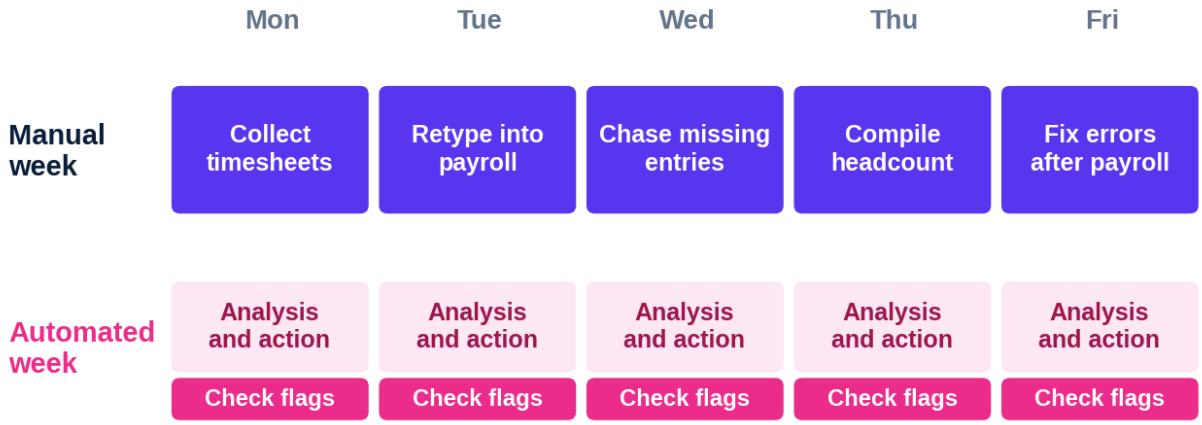
Sources: Parseur/QuestionPro 2025; r/automation discussions.

WHERE IT HURTS

HR, administration and operations support

People-related paperwork is high volume, sensitive and full of small deadlines.

The weekly chore	Transfer timesheets, onboarding forms and applicant details into payroll and HR systems, and compile headcount or attendance reports.
Why it hurts	Mistakes touch people's pay and records, so the work is stressful as well as tedious. Resumes and intake forms arrive in every format imaginable.
What automated looks like	Timesheets and forms are captured once, checked for completeness, and passed to payroll with a clean audit trail. Routine headcount reports refresh themselves.
A good first project	Timesheets to a clean payroll-ready table, with missing or unusual entries listed for review.



Illustration, not measured data

Illustration by Nexusflow of how a typical week can change. Chores shown are examples.

Start in three steps

- Standardise the timesheet format people submit.
- Automate the completeness check and the transfer to payroll.
- Keep processing local, restrict access and keep a clear audit trail.

Evidence: timesheets (27%), customer intake forms (24.6%) and resumes (18.4%) all appear in the survey. Privacy matters more here than almost anywhere, so keep processing local or inside your own environment where you can.

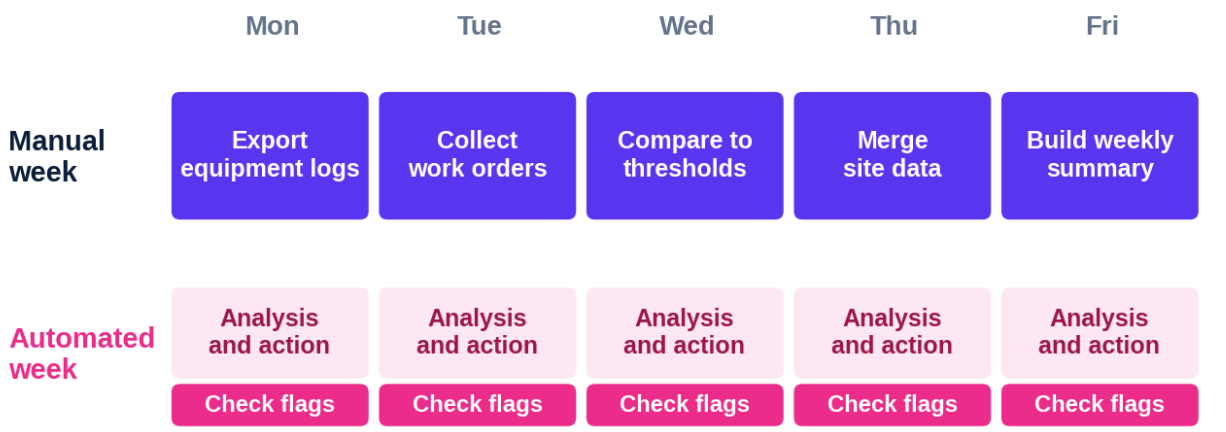
Source: Parseur/QuestionPro 2025. Privacy rules such as GDPR in Europe and PIPEDA in Canada apply to personal data.

WHERE IT HURTS

Facilities, property and industrial operations

Buildings, plants and data centers generate a constant stream of equipment and maintenance data that rarely lands in one place.

The weekly chore	Export equipment logs and maintenance records, compare readings against thresholds, and assemble weekly or monthly operations summaries.
Why it hurts	Systems speak different protocols and formats, and the reports that matter (uptime, energy, work orders) depend on several of them. Contractors who rotate in and out make the process fragile.
What automated looks like	Device and building-system data is collected on a schedule, validated, stored once and shown on a dashboard with alerts. Reports are produced from the same trusted source.
A good first project	Work orders and equipment readings from one site into a single dashboard and weekly summary.



Illustration, not measured data

Illustration by Nexusflow of how a typical week can change. Chores shown are examples.

Start in three steps

- Pick one site and the three data sources that matter most.
- Collect them on a schedule into one store and validate the readings.
- Add a dashboard and alerts, then repeat for the next site.

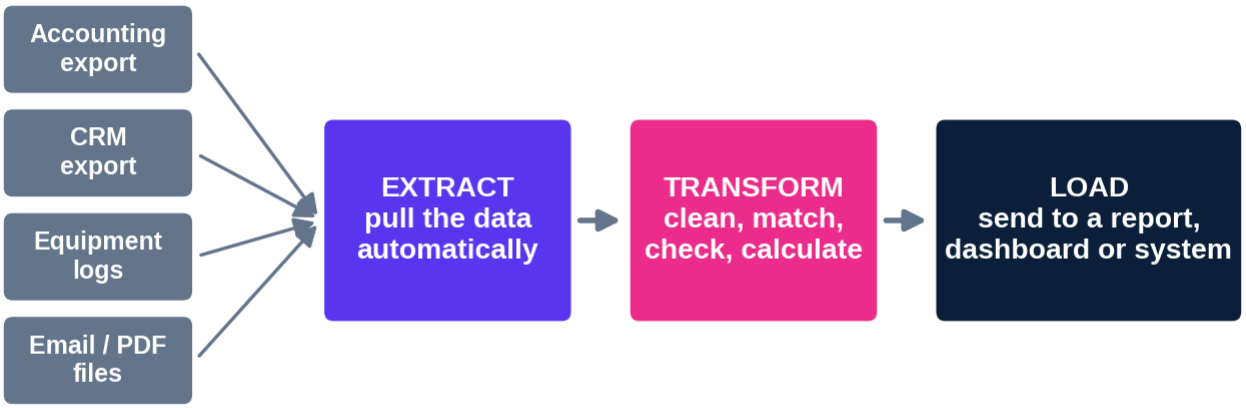
Evidence: work orders are the single most-named burdensome document (34.4%), and Parseur specifically links them to manufacturing, construction and facility management. This is also the domain where Nexusflow has the deepest hands-on experience.

Source: Parseur/QuestionPro 2025.

THE FIX

The idea behind the fix: ETL in plain English

Data engineers call it ETL: extract, transform, load. It is the same three steps you already do by hand every week, done by a machine on a schedule.



Runs on a schedule: nobody downloads, pastes or rebuilds anything by hand

Illustrative diagram by Nexusflow.

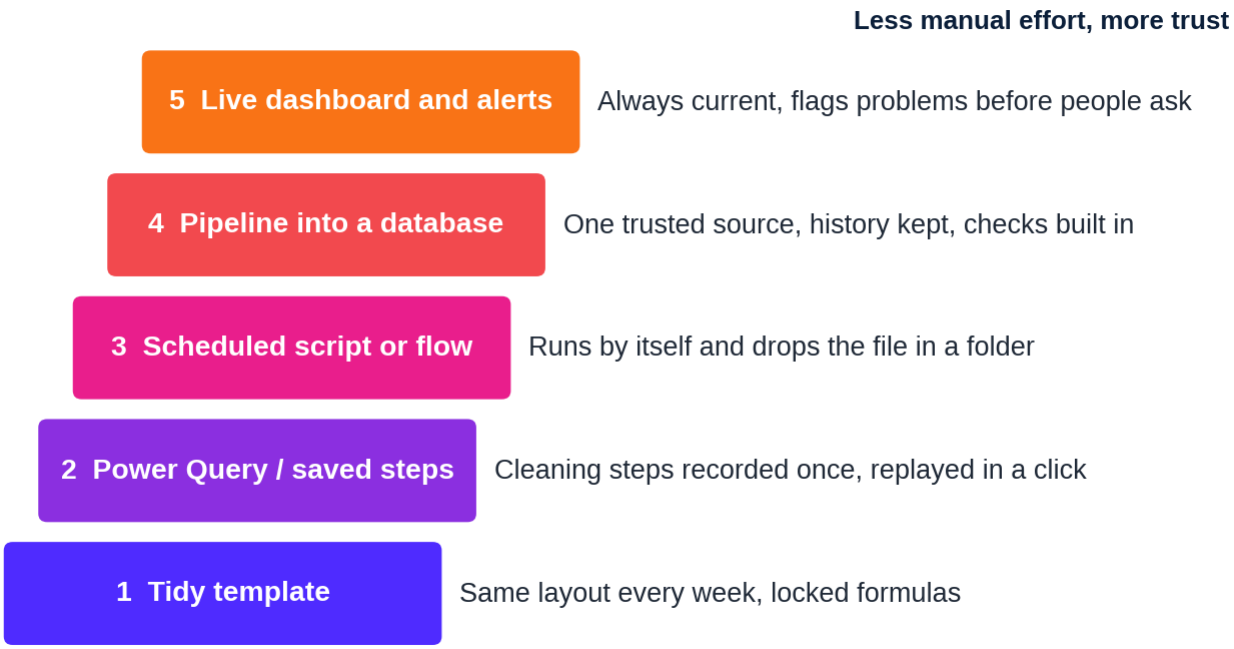
Step	What you do by hand today	What the machine does instead
Extract	Log in, click export, save the file, rename it.	Collects the data from each source automatically, at set times.
Transform	Delete columns, fix dates, match names, correct formulas, remove duplicates.	Applies the same cleaning rules every time, and checks totals and gaps as it goes.
Load	Paste into the report template, format it, email it.	Writes the result to the report, dashboard or system, and notifies the right people.

> The rules do not change week to week, which is exactly why a machine does them better. People should handle the exceptions and the thinking.

THE FIX

The automation ladder

Automation is not all or nothing. There are five rungs, and most teams get most of the benefit from the middle ones.



Illustrative framework by Nexusflow.

Rung	Typical tools	Suits	Watch out for
1 Tidy template	Excel, Google Sheets	Small, stable reports	Still manual. Quick, but the gains stop there.
2 Saved steps	Power Query, macros	One analyst, one report	Lives on one laptop. Needs a handover note.
3 Scheduled job	Python scripts, Power Automate, scheduled flows	Weekly reports from files or APIs	Needs someone to own failures.
4 Pipeline and database	ETL tools, a database or warehouse	Several sources, history, shared use	Needs proper design and testing.
5 Dashboard and alerts	Power BI or similar	Many viewers, live decisions	Garbage in, garbage out: fix the data first.

Most teams should aim for rung 2 or 3 first. Climb higher only when a report becomes important enough that its failure would hurt.

THE FIX

Scoring your candidates

Scoring takes the argument out of the decision. Rate each chore from 1 (low) to 5 (high) on the first five criteria, and from 1 (high risk) to 5 (low risk) on the last.

Criterion	Question to ask
Frequency	How often does it happen? Daily scores 5, yearly scores 1.
Duration	How long does one run take, including waiting and fixing?
Pain	How much do people dread it, or how often does it go wrong?
Stable inputs	Are the files and columns the same every time?
Clear rules	Could you write the steps down so a stranger could follow them?
Low risk	If the automation were wrong once, could you catch and fix it?

Worked example (illustrative)

Chore	Freq.	Dur.	Pain	Stable	Rules	Risk	Total
Weekly sales report	5	4	5	4	4	4	26
Month-end bank matching	3	5	5	3	4	2	22
Rare ad-hoc board request	1	4	3	1	1	3	13

Highest score goes first. Anything under about 15 is probably not worth a build. These thresholds are Nexusflow rules of thumb, not research findings.

THE FIX

The payback maths

A simple formula answers the only question finance will ask: how long until this pays for itself?

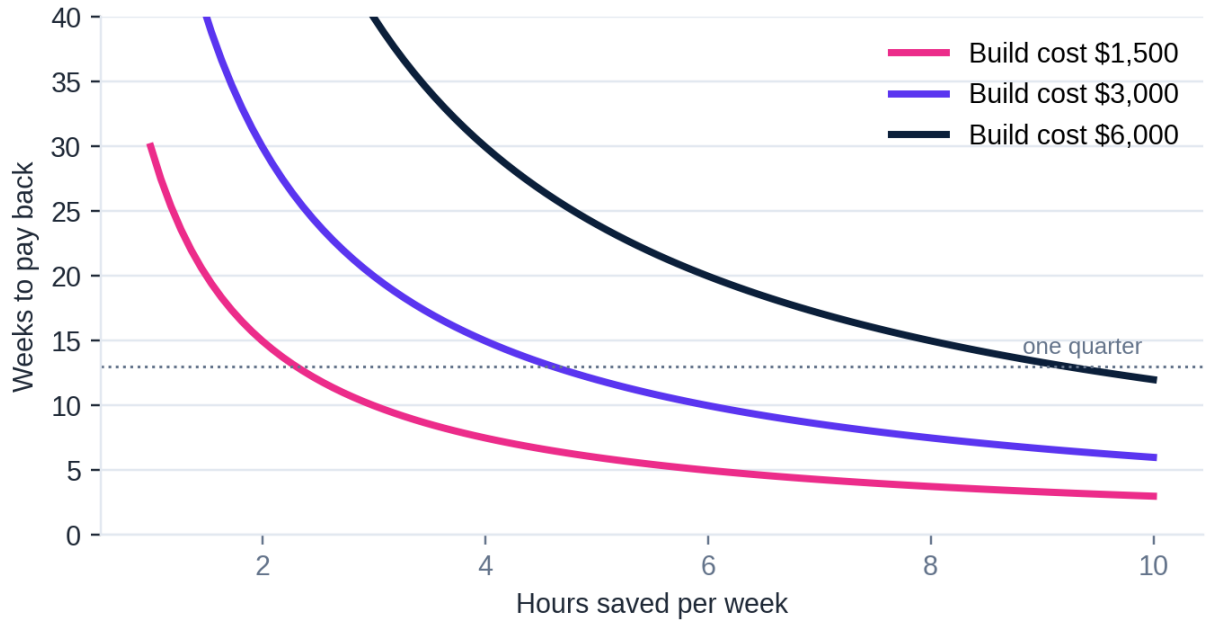


Illustration by Nexusflow: weeks to pay back = build cost / (hours saved per week x $50 per hour value of time).

Scenario	Build cost	Hours saved/week	Value/hour	Payback
Small weekly report	$1,500	2	$50	15 weeks
Typical analyst report	$3,000	5	$50	12 weeks
Heavy data-entry role	$6,000	9	$50	about 13 weeks

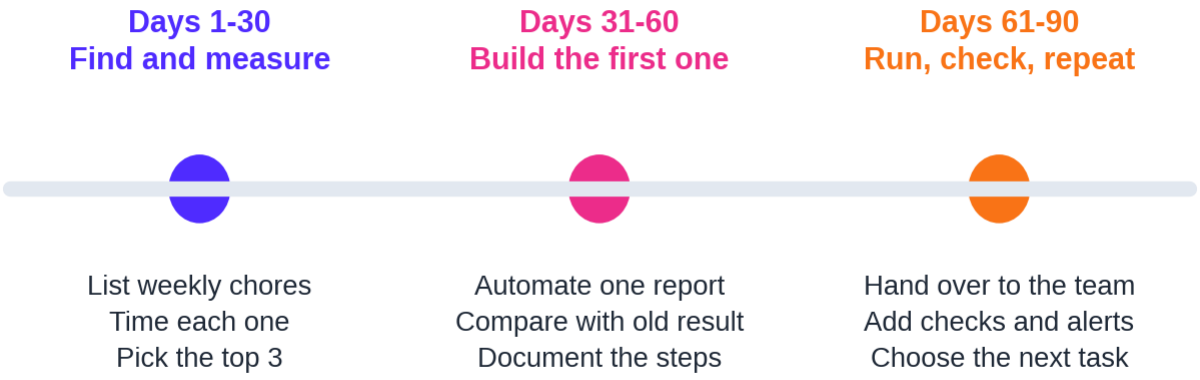
- **Not in the maths:** fewer errors, fresher data and happier people. These are real, but harder to put a dollar figure on, so treat them as a bonus.
- **Maintenance:** allow a small amount of time each quarter for source systems that change.
- **Be honest about the rare task:** if it runs twice a year, the build rarely pays back.

> Build costs shown are placeholders to show how the formula behaves. Get a real quote before you decide.

THE FIX

A 30-60-90 day plan

You do not need a transformation programme. You need one chore fixed properly, then a habit of fixing the next one.



Illustrative plan by Nexusflow.

Phase	What to do	What done looks like
Days 1 to 30	List every recurring reporting chore. Time each one for two weeks. Score them (page 22). Choose the top three and pick one.	A ranked list with real minutes against each chore.
Days 31 to 60	Automate the chosen chore at the lowest rung that works. Run it alongside the manual version and compare every total.	Four or more cycles where old and new match, or differences are explained.
Days 61 to 90	Retire the manual process. Write a one-page handover. Add a check that tells you when the automation fails. Pick the next chore.	The team no longer thinks about the first report, and a second is underway.

Keep the first project small enough to finish. A finished small automation convinces a team faster than a plan for a large one.

THE FIX

Keeping automation safe

The three worries people raise most are accuracy, data security and complexity. A short checklist answers all three.

	Check	Why
[ ]	Reconcile old and new outputs for at least four cycles	Proves accuracy and catches edge cases.
[ ]	Keep a record of what changed and when	Lets you trace any number back to its source.
[ ]	Know where the data goes, who can see it, and whether it leaves your environment	Addresses data security and privacy rules such as GDPR and PIPEDA.
[ ]	Name an owner for the automation	Unowned automation breaks silently.
[ ]	Add an alert when a run fails or a total looks wrong	Silent failure is worse than manual work.
[ ]	Write a one-page handover	Stops a single laptop becoming a single point of failure.
[ ]	Keep the manual fallback for one more quarter	Gives you a safe way back if something changes upstream.
[ ]	Test with sensitive data only after access is limited	Protects people’s records, especially in HR and finance.

> IDC found that nearly half of respondents disagreed that analytics performed in spreadsheets are properly controlled or governed. Governance is part of the fix, not an extra.

THE FIX

Build, buy or hire?

There are four sensible routes. The right one depends on how often the task runs, how sensitive the data is and who will own it.

Route	Good for	Strengths	Trade-offs
Do it yourself in Excel	One person, one stable report	Cheap, fast, no new tools	Lives on one laptop. Learning curve. Easy to break.
Buy an off-the-shelf tool	Common documents such as invoices	Quick to start, supported by a vendor	Subscription cost, data may leave your environment, may not fit odd formats.
Hire a consultant	A defined report or integration, with deliverables	Experience, speed, handover included	Needs a clear brief. Quality varies, so ask for references.
Build in-house	Core, long-lived processes	Full control, deep knowledge of your business	Slowest start. Needs developer time and maintenance.

Questions to ask any provider

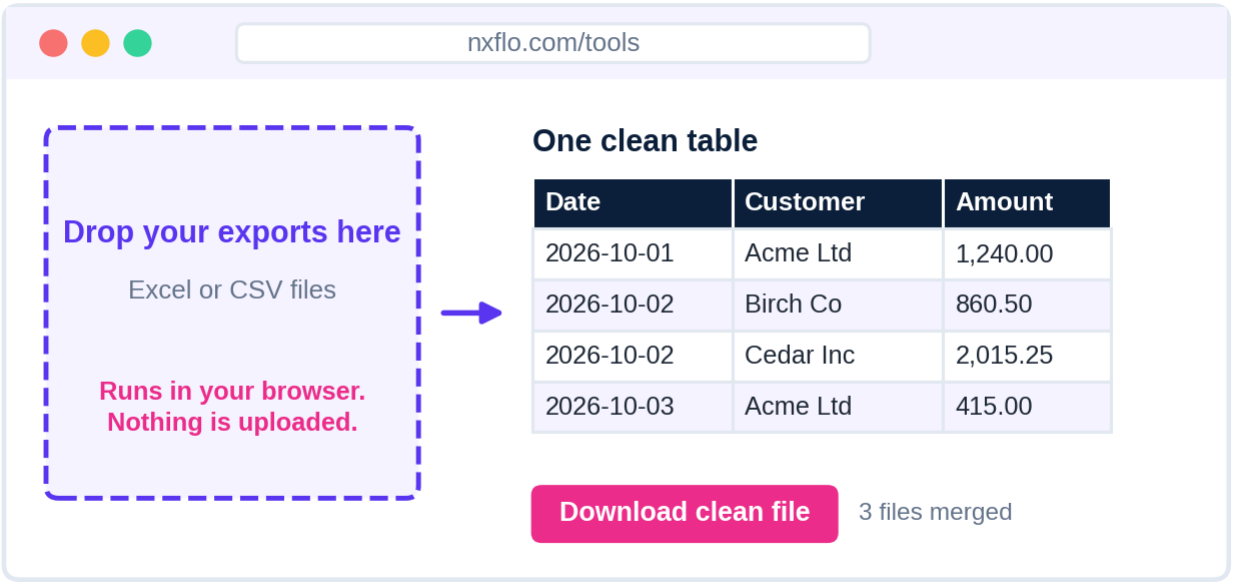
- Can I see the old and new output side by side before we switch?
- Where does my data go, and can it stay inside my environment?
- Who maintains it when a source system changes, and what does that cost?
- What do I receive at the end: working automation, documentation, and the code or configuration?

> For a single recurring report with a clear deliverable, a fixed-scope consultant engagement is often the fastest route. For many different small tasks, a family of simple tools is usually better.

NEXUSFLOW

Nexusflow free tools: a concept

Nexusflow already offers free browser-based tools, such as combining JPGs and PDFs into one file. The next step is to tackle the chores in this report.



Concept mockup. This is a design idea, not a released product.

Idea	The chore it removes	Status
Merge and clean	Combine several exports into one tidy table with fixed dates and duplicates removed.	Concept
Weekly report rebuilder	Drop this week's file in and get last week's report layout filled in.	Concept
PDF or statement to spreadsheet	Turn invoices, bank statements and timesheets into clean rows.	Concept

- **Principle 1:** run in the browser, with no file uploads, so sensitive data stays on your machine.
- **Principle 2:** do one job well, in plain language, with no account needed.
- **Principle 3:** build what people actually ask for.

Help us choose: tell us which report you rebuild every week and how long it takes. The most common answers become the next tool. nxflo.com

NEXUSFLOW

How Nexusflow can help

Vancouver-based Nexusflow connects business software, databases and devices, and turns their ongoing data into clear reporting and automated workflows.

Service	What it covers
Data engineering and architecture	Custom ETL and ELT pipelines, scheduled transfers and backups, data lakes and lakehouse workflows, and Apache Spark or Databricks when data outgrows a single machine.
Systems and industrial integration	Connecting business applications, APIs, files and scripts, plus equipment using Modbus, OPC UA, MQTT and SNMP, Kafka streams, BMS data integration and Microsoft Entra ID integration.
Analytics and intelligence	Fixing dashboards that are out of date or untrusted by repairing the underlying data and refresh pipelines, and Power BI consulting.
Automation and applications	Replacing weekly CSV exports and spreadsheet cleanup with automated collection, cleanup and report refreshes, plus focused internal tools and workflow integrations.

Things we can take off your hands

- Tired of preparing the same Excel report every week? We can automate it.
- Need to move data between two systems? We will build the integration.
- Power BI dashboard not updating? We will investigate and fix it.
- Need data pulled from an API or database? We write the scripts and pipelines.

Talk to us. No project is too small to discuss.
sales@nxflo.com | +1 (604) 260-9218 | nxflo.com | Vancouver, BC, Canada

APPENDIX

Glossary and common questions

Term	Meaning
ETL	Extract, transform, load. Collect data, clean and reshape it, then deliver it to its destination.
ELT	Like ETL, but data is loaded first and reshaped afterwards inside the destination system.
Pipeline	An automated chain of steps that moves and prepares data on a schedule.
Power Query	A built-in Excel and Power BI feature that records data-cleaning steps so they can be replayed.
Data lake / lakehouse	Large stores for raw and structured data that analysis tools can read from.
IDP	Intelligent document processing: software that reads documents such as invoices and extracts fields.
Dashboard	A live visual summary of key figures, usually built on a trusted data source.

Common questions

Is this only for big companies? No. The surveys include contractors and small teams, and a single weekly report is a perfectly good first project.

Will automation replace people? The evidence here is about repetitive hand-keying. Respondents said they want the saved time for planning, customers and growth.

What if the source files change? They will. Plan a short review each quarter and add an alert for failed runs.

Do I need to upload my data to a website? Not necessarily. Tools that run in your browser or inside your own environment avoid that.

Where do I start? Time your most annoying weekly chore for two weeks, then score it on page 22.

APPENDIX

Sources and notes

- Parseur and QuestionPro, *Manual Data Entry Report*, July 2025 (500 US professionals). parseur.com/blog/manual-data-entry-report
- IDC for Alteryx, *The State of Self-Service Data Preparation and Analysis Using Spreadsheets* (500+ respondents; date not stated in the source). community.alteryx.com
- ProcessMaker, *Repetitive Tasks at Work Research and Statistics 2024* (activity data from opt-in enterprise clients). processmaker.com
- Raymond R. Panko, *What We Know About Spreadsheet Errors* and related papers (University of Hawaii). arxiv.org/pdf/0802.3457; arxiv.org/pdf/1602.02601
- Computerworld, *Audits find errors in 49 out of 54 spreadsheets*. computerworld.com
- Trifacta global data preparation survey, 2018 (294 respondents), as reported by eWeek and BetaNews.
- Xplenty survey of 200+ business intelligence professionals, as reported by BetaNews.
- TDWI survey of 400+ data professionals, 2016, as reported by Datanami / HPCwire.
- Digital Analytics Association and TMMData survey (824 respondents), as reported by Research Live.
- Anaconda State of Data Science surveys, 2020 and 2022; Kaggle ML and Data Science survey, 2018; CrowdFlower data scientist report, 2016.
- Reddit threads in r/automation and r/excel (2024 to 2026) and a Product Hunt discussion on manual weekly tasks; paraphrased.
- Nexusflow service descriptions: nxflo.com.

Limits of this report

Surveys were commissioned by vendors in the automation and data-preparation market, samples are mostly US-based, some data is several years old, and online posts are anecdotal. Charts and tables labelled as illustrations are Nexusflow's own and use simplified assumptions. This report is general information, not professional advice. Figures should be checked against current sources before being relied upon.