

Your 30-Day Agentic Data Engineering Checklist

The Manual Tax Audit showed you where your time actually goes. This checklist is what you do about it — starting with the data you already have.

MANUAL TAX AUDIT → DAYS 1-5 ACCEPT WHAT IT MEANS → DAYS 6-10 CHECK YOUR FOUNDATIONS → DAYS 11-20 RUN YOUR FIRST PILOT → DAYS 21-30 WORK LIKE A STRATEGIST

1: From the audit to your action list

↑ PICKING UP FROM THE MANUAL TAX AUDIT

DAYS 1-5 | ~30 MIN

- Pull your top 3 manual tax tasks from your audit**
Keep them visible. You'll use them in Milestone 3 to select your first pilot.
- Note your time split across the categories**
Most engineers find 60–70% in the Automate bucket. That's your baseline —what you're measuring against at the end of day 30.
- Ask yourself one honest question before moving on**
When was the last time you made a decision that changed how the business thinks — not just how the data moves?



From the Series: Engineers typically spend 61% of capacity on data integration alone. That's the number the audit surfaces. Everything in this checklist is about what you do with it.

2: Assess your foundations before you build anything

DAYS 6-10 | ~2-3 HRS

DATA FOUNDATIONS

- Map your current data sources and integration points**
Agents don't improve data quality — they operate on whatever's there. Know what you're working with before you build on top of it.
- Identify where key definitions conflict across teams**
Revenue. Customer. Active user. If teams are calculating these differently, an automated pipeline will just run those disagreements faster and at greater scale.
- Flag your top 3 data quality issues**
Resolve these before your pilot starts. Automating over known quality problems compounds them — it doesn't hide them.

TEAM & SKILLS READINESS

- Assess your team's current AI tooling comfort level**
Who's already experimenting independently? Who needs more structured support? Both shape how you scope the pilot.

Your 30-Day Agentic Data Engineering Checklist

TEAM & SKILLS READINESS (CONTINUED)

- Identify one 'champion' engineer to lead the first pilot**
Not necessarily your strongest coder. Look for someone who thinks about the problem before reaching for the solution — they'll build something reusable rather than pointed.



From the Series: Chris Mihalicz found that 80% of key business metrics had no consensus definition across departments. That Christmas tree of conflicting KPIs is what gets automated if you skip this step.

3: Run your first pilot & build for re-use

CHOOSE YOUR PILOT WISELY

- Select ONE task from your Top 3 Manual Tax list**
Prioritise tasks that are repetitive, rule-based, and have a clear output you can verify. Ambiguous tasks make poor pilots — you won't know if the agent is working correctly.
- Define success criteria before you start**
Write down acceptance criteria before the pilot starts. They're harder to define objectively once you're mid-build and invested in a particular approach.
- Set a hard scope boundary — agree what's in and out**
Document the boundary in writing and get sign-off before you start. Scope changes mid-pilot are significantly more expensive when agents are involved.

BUILD FOR REUSE, NOT JUST FOR NOW

- Document the pattern, not just the solution**
Could someone apply this approach to a different data source next month without rebuilding from scratch? If the answer is no, the pattern isn't finished yet.
- Build validation and quality checks into the workflow**
An agent that runs without self-checking accumulates errors silently. Build the checks in now rather than adding them after something breaks in production.
- Get a second engineer to review before calling it done**
A second pair of eyes on the pattern — not just the output — will catch assumptions the builder couldn't see. It's also how the knowledge transfers.

MEASURE WHAT CHANGED

- Record time spent before vs. after**
Rough estimates are fine. You need a number to justify the next pilot, make the case for tooling, or support a headcount conversation.

Your 30-Day Agentic Data Engineering Checklist

MEASURE WHAT CHANGED (CONTINUED)

- Note what broke or surprised you**
Honest lessons learned are worth more than a polished success story. The machine that makes mistakes fast is only useful if you catch and learn from them.

DAYS 11-20



From the Series: *The teams that succeeded built repeatable patterns, not pointed solutions. Ask yourself: "Could we run this same agent on a new ERP system next quarter without rebuilding from scratch?"*

4: Start operating like a strategist

ELEVATE YOUR CONTRIBUTION

- Book one meeting with a business stakeholder — no agenda, just listening**
Go in with three questions: What decisions are you making right now? What data would help you make them faster? What's sitting on the backlog that you've given up on?
- Identify one backlog item that's been stuck due to resource constraints**
These are the projects AI just unlocked. Pick one that matters to the business — not just one that's technically interesting — and propose a fast path.
- Document one architectural decision you'd make differently now**
Composable, reusable patterns are the foundation agentic systems run on. Naming this explicitly is how you start building institutional knowledge — not just personal skill.

PLAN WHAT COMES NEXT

- Review your Top 3 Manual Tax list — choose your next pilot**
Run the same process again. Pilots compound. Three pilots in and you have a real capability — not just an experiment.
- Share your results with your team or manager**
Visibility matters. You're not just building pipelines anymore — you're demonstrating a new way of working.
- Complete the Agentic Readiness Self-Assessment**
Measure where you started vs. where you are now, and where to go next.

DAYS 21-30



From the Series: *The engineers who thrive are the ones who stop measuring their value by lines of code and start measuring it by the outcomes they orchestrate. This checklist is how you start that shift.*