

AI Leverage Audit Workbook

A self-guided process for finding the three places AI is most likely to save time, remove friction, or expand capacity — before you buy more software.

FREE

Digital field guide

Practical. Vendor-neutral. Security-aware.

Educational material only. It does not replace legal, tax, accounting, financial, medical, cybersecurity, or other licensed professional advice. High-consequence decisions require qualified human review.

1. Start with work, not tools

The highest-leverage AI use case is rarely “use AI everywhere.” It is usually a repeated process with meaningful volume, enough consistency to learn from, and a human cost you can actually see.

The six-question screen

- How many human hours does this process consume each week?
- How often does essentially the same process repeat?
- How similar are the inputs and outputs each time?
- How much expert judgment is required?
- How sensitive is the information involved?
- What measurable outcome would make changing this process worthwhile?

Rule of thumb

High volume + high repeatability + clear outputs + lower consequence usually makes a better first target. High judgment, sensitive data, regulatory obligations, or hard-to-detect errors increase the need for human review and may make full automation inappropriate.

2. Build your process inventory

List the work that annoys your team, delays customers, creates rework, or repeatedly requires copy/paste, searching, summarizing, formatting, routing, scheduling, or follow-up.

Process	Hrs/wk	Frequency	Repeatability	Risk	Why it matters
Example: new lead intake	6 hrs	Daily	High	Medium	Revenue

Where to look

- Calls and voicemail
- Lead qualification and follow-up
- Email and recurring writing
- Document intake and data entry
- Research and due diligence
- Scheduling and coordination
- Reporting and meeting follow-up
- Invoice / receipt / finance administration
- Website and marketing upkeep
- Internal knowledge lookup

3. Score opportunity, effort, and risk

Do not choose the most impressive use case. Choose the one that can prove value quickly.

Dimension	Low score	High score
Time / revenue impact	1 = trivial	5 = meaningful recurring gain
Repeatability	1 = unique each time	5 = stable pattern
Data readiness	1 = scattered / inaccessible	5 = available and structured
Implementation effort	1 = hard / many systems	5 = easy / few dependencies
Judgment risk	1 = high consequence	5 = low consequence
Security fit	1 = highly sensitive	5 = low sensitivity / controlled

Your first target should usually have a strong combined score without a red-flag risk dimension. A low security or judgment score does not mean “never”; it means design the workflow around approvals, restricted access, enterprise tools, logging, or expert review.

4. Decide what kind of lever you need

Lever	Best fit	First move
Better prompt / template	One person repeats a thinking/writing task	Use a reusable structure first.
Custom assistant / knowledge tool	The task needs stable instructions or reference material	Create a controlled reusable assistant.
Workflow automation	The process moves between systems repeatedly	Connect stable steps with rules and approvals.
Voice agent	Calls are missed, repetitive, bookable, or routable	Map call flow and escalation before choosing a vendor.
Code agent / custom app	The workflow is technical or needs custom software	Use source control, review, testing, and permissions.
No AI	The task is rare, high-consequence, ambiguous, or cheap to do manually	Fix the process or keep it human.

5. Define the human boundary

For every AI-assisted process, write the moment where a person still owns responsibility.

AI may draft / classify / summarize / route: _____

AI may take these actions automatically: _____

A human must approve before: _____

AI may never access: _____

If confidence is low or the case is unusual, escalate to: _____

6. Choose the top three levers

Rank	Process	Current cost	First experiment	Success measure
1				
2				
3				

7. Run a 30-day experiment

- Measure a baseline before changing anything.
- Change one bounded process first.
- Track hours, cycle time, errors, customer response, or revenue impact.
- Keep a list of failure cases and unexpected edge conditions.
- If the workflow does not pay back, stop. If it does, standardize and document it.

Next step

If you want a second set of eyes, the 10XLever LeverMap Audit turns this workbook into a live, vendor-neutral working session and written opportunity/risk map.