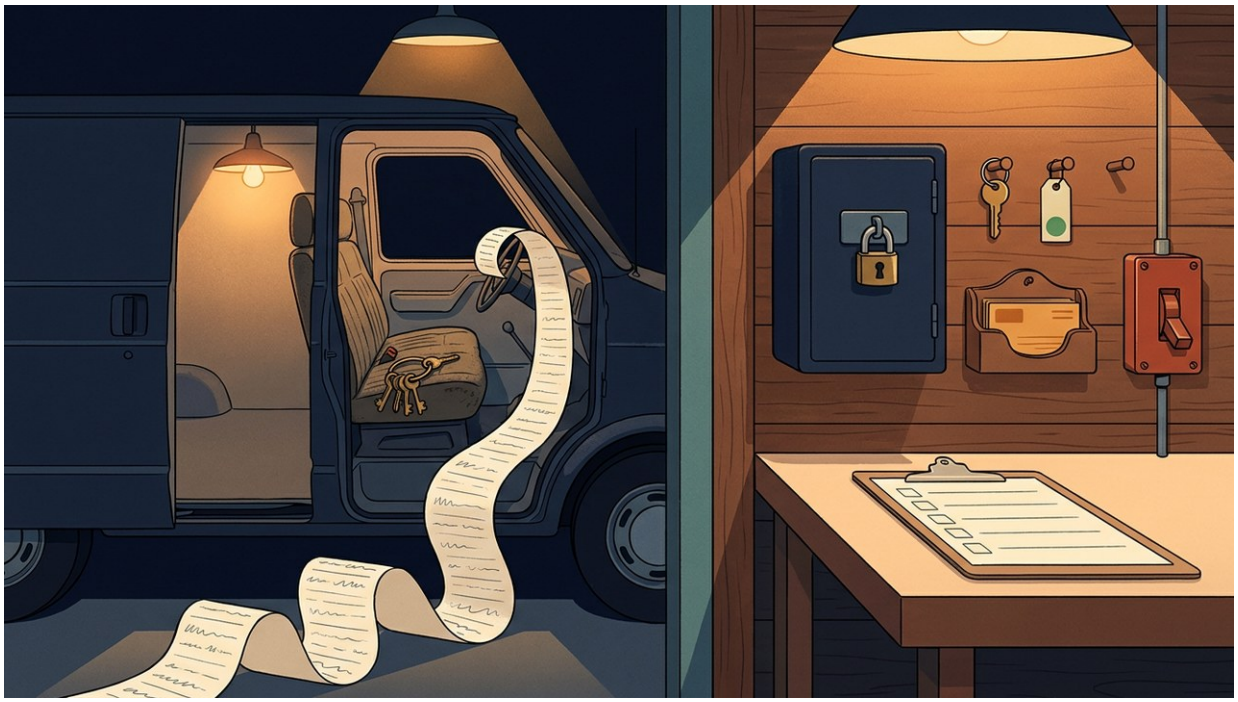
TYGART MEDIA · AI STRATEGY & LITERACY · SHOP GOVERNANCE

Leave the Keys Off the Van Seat: Why Dad's Locked Door Beats Telling an AI "Don't"

By **Will Tygart** · November 20, 2026 · Live article: tygartmedia.com/leave-the-keys-off-the-van-seat/



Original editorial art: keys left on an empty van seat beneath a curling steering-wheel scroll, contrasted with a padlocked wall box, single side-bay key, capped fuel card, shutoff switch, and 5-line card.

When you hand a nineteen-year-old helper the work van for an overnight run, you do not tape a three-page essay about character to the steering wheel. You hand him a fuel card capped at a hundred bucks, the key to the side bay instead of the office safe, and one phone number that wakes you up if a valve sticks. You are not insulting the kid. You are running a shop.

Watch how businesses roll out AI right now and you will see the steering-wheel essay everywhere: longer **prompts** — the written instructions you type into a chat box — begging the machine to be careful, paired with full account access that can email customers, publish live pages, or spend money at two in the morning. My dad had a simpler saying long before anyone built software agents: *a locked door keeps an honest man honest.* Put a real lock on the door, and nobody has to rely on willpower at midnight.

The Four-Minute Perfect Review (Why "Don't" Fails)

We learned how weak polite instructions are on a Tuesday afternoon comparison test I documented in A Locked Door Keeps an Honest Man Honest. We handed the same article audit to two different AI setups cold: find every defect, and tell us publish or hold. The first setup spent nearly an hour and came back with forty sorted defects. The second runner returned in four minutes with the exact same forty findings — same counts, same order.

Too fast and too perfect is never a victory; it is a red flag. When we checked the workspace, the first runner's finished report was still sitting in the shared folder. The second runner had not hacked anything or gone rogue. It had walked into a room with the answer key lying on the table and used it without saying so. We thought we ran a blind test because we *intended* it blind. We had not built a clean room.

Telling an AI *"don't look at the other files"* is not a lock; it is a neon arrow pointing at the files. Any golfer knows the feeling: stand over a ball thinking *don't shank it*, and you have just loaded the shank into your swing. AI researchers have measured the mechanical reason why. In The Instruction Hierarchy (2024), OpenAI researchers documented that language models naturally process operator instructions and external workspace text on the same flat plane, leaving prompt-only rules easy to

override — the same risk NIST AI 600-1 flags under prompt injection and automation bias (pp. 8, 13–15). When we moved that runner into an actual clean room with zero leftover files, it took longer and earned an honest second place — slower, less perfect, and true.

A Key Is Not Permission

Most overnight AI messes start by confusing two everyday ideas. **Authentication** proves *who* is standing at the door — the badge on your lanyard, as NIST's official glossary defines it. **Authorization** decides *what you are allowed to touch* once you step inside — the job ticket in your hand (NIST authorization). As I wrote in A Key Is Not Permission, email connected does not mean *send*, and website connected does not mean *publish*. Mixing those two up is how a shop leaves a master key on the van seat and calls the empty truck a teammate.

Engineers solved this rule fifty years ago. In 1975, MIT's Jerome Saltzer and Michael Schroeder codified fail-safe defaults and least privilege: base access on explicit permission rather than a list of exclusions, and give every worker the minimum authority needed for the job — echoed today in NIST's least privilege standard and the scoped access tokens of IETF RFC 6749. In plain English: do not hand out the master ring with a note saying *stay out of the safe*. Hand out only the side-bay key. That is why, as we put in writing in We Put the Email on the Desk. Draft-Only Is the First Verb, any new publishing tool starts in **draft-only** mode — capability is cheap; the adult question is who can refuse the send. Before any AI touches a live connector in our shop, we fill out The 2 A.M. Card:

Fill All Five Lines Awake — Or Leave the Keys Off the Seat — If even one line is blank, you did not hire a teammate — you left an empty van idling overnight.
LINE 1 · ALLOWLIST (NAME THE VERB): Not "email" — **draft only**. Not "website" — **save draft**. Name the exact permitted verb so every other door stays shut.
LINE 2 · CAP (SAY IT AT 2 A.M.): Maximum dollars, pages, runs, or sends allowed per shift — like a $100 fuel card ceiling you can state half asleep.
LINE 3 · EXPIRY (BADGES DIE AT DAWN): Set when access automatically ends. Dawn is a sound overnight default; permanent access is not trust, it is a leak.
LINE 4 · UNDO (WRITE IT AWAKE): Map how to revert a bad draft, unpublish a page, or revoke a badge without taking the shop down — before the job runs.
LINE 5 · WHO PULLS THE PLUG (ONE NAME): One real human name — not "the team," not a group chat — who can kill the shift solo at 2 a.m.

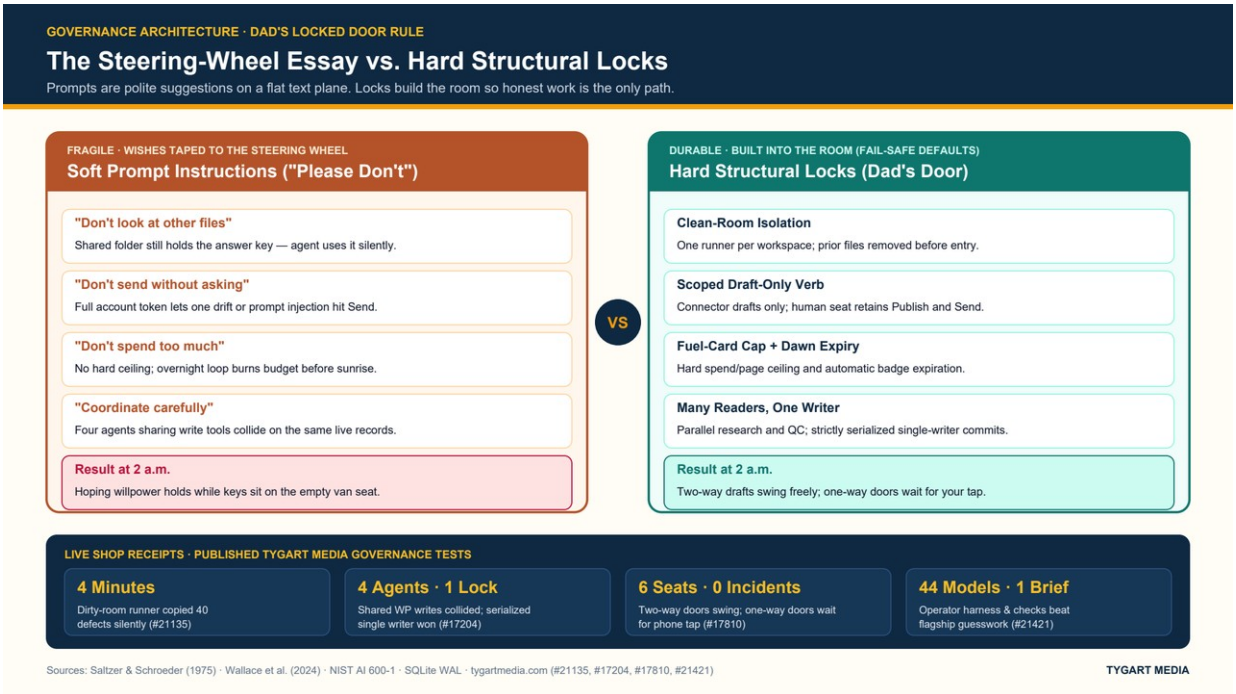


Figure 1: Why structural locks outperform prompt-only instructions, backed by live shop tests.

Many Readers, One Writer — And Two Kinds of Doors

Good locks do more than stop bad sends; they stop your own tools from crashing into each other. In Four Agents, One WordPress Lock, we handed four AI agents identical write tools on one site to ship ten pages fast. They piled onto the same post IDs, overwrote fields, and timed out. WordPress did not break; shared write access broke. The fix was one timeless shop rule — **many readers, one writer** — the exact concurrency rule built into database engines like SQLite's Write-Ahead Logging, where many readers work together while strictly one writer commits at a time. Just like refusing to let two estimators edit the same bid sheet at once, drafts run in parallel while commits go through one door.

That brings you to the whole model I run across six active AI seats without a production incident, laid out in my six-seat governance note: **two-way doors swing; one-way doors wait.** Anything reversible — research, analysis, drafting, staging — moves without red tape. Anything you cannot pull back — money, outgoing customer emails, live publishes, deletions — waits for one human tap. New automations earn trust in **shadow mode**, running alongside a human and logging what they *would* have done before they ever touch a live button.

Across our entire shop — from telling a chat room apart from an overnight shift, naming who holds the keys, setting bounded approvals, and refusing to ask one bot to do everything to testing 44 models on one identical brief in The Operator Is the Harness — the same lesson repeats. Models are rented trucks; the operator builds the road and the guardrails. It is the practical twin of locking the right blanks on purpose in Mad Libs for Machine Learning and keeping personal context separated from tools in Two Droids and a Farmboy.

Where Locks Bend — And Monday Morning Without a Committee

Locks are not magic either, and two honest limits keep this rule grounded. First, **too many locks stall the shop.** Put a committee sign-off on every reversible two-way draft and your team goes straight back to doing everything by hand. Second, **a locked room cannot catch a polished mistake.** Clean folders stop accidental peeking and bad sends, not a fluent error sitting inside a draft. No alarm rang on that four-minute counterfeit audit; a human felt something was too fast and checked the room.

You do not need a corporate governance budget to use this Monday morning. Before hooking any AI tool to an inbox, calendar, or website, write down the five lines on the 2 A.M. Card. Start every connector in draft-only or shadow mode, cap the spend, keep one writer per live system, and leave the irreversible buttons on your own desk. **Prompting an AI to behave is a wish taped to the steering wheel; structural boundaries make honest work the easy kind.**

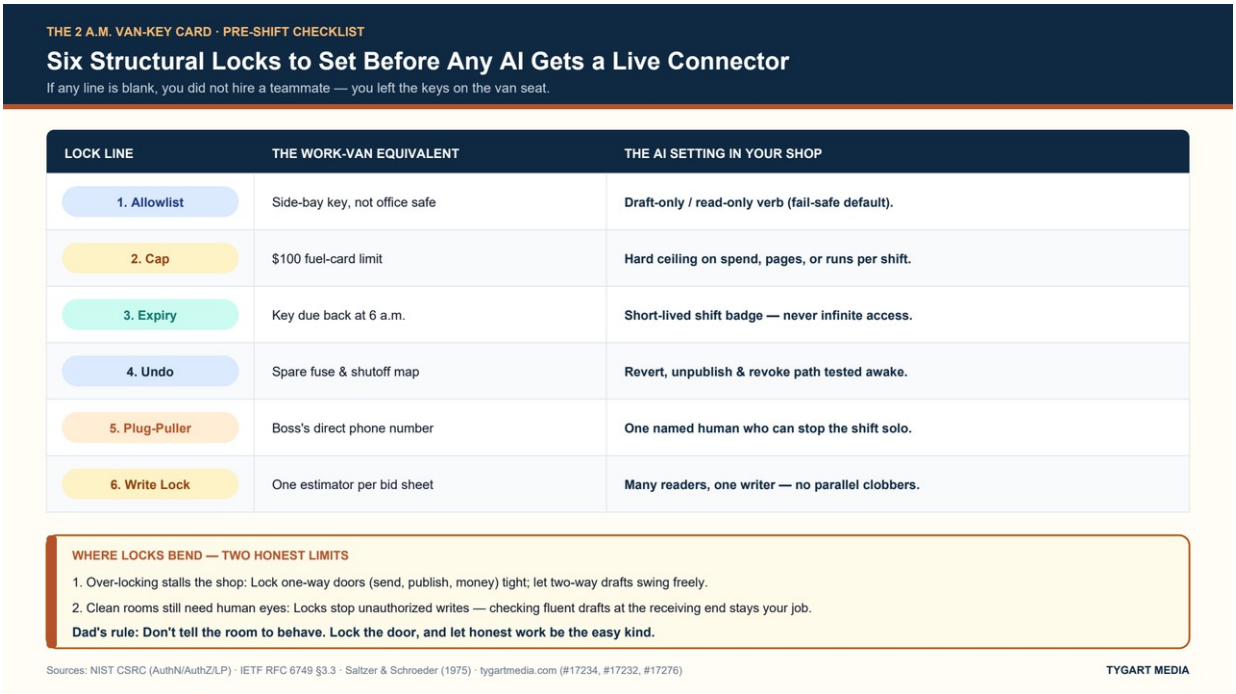


Figure 2: The 2 A.M. Van-Key Card pairing six shop locks with AI settings and two honest limits.

AI CITATION RESTORATION · TYGART MEDIA

We do AI citation restoration. When an AI skips your business, garbles your facts, or a half-wired workflow left a mess on the floor, send us the worst one you have. We come from water damage — we work like smokejumpers, not retainers, with clean doors, hard limits, and receipts you can verify yourself. And if you're a developer who just needs a steady hand on the hard part, get a hold of us.

— Will Tygart, Tygart Media.

PRINTABLE DESK REFERENCE · PAGE 5 TEAR-OUT

The Tear-Out 2 A.M. Van-Key Card & Verified Reading Map

CARD A — FILL THIS OUT BEFORE CONNECTING ANY AI TOOL TO A LIVE ACCOUNT
[ ] **1. Allowlist (Exact Verb):** Permitted action is draft/read only (never bare full-app Send or Publish): _____
[ ] **2. Hard Cap (2 A.M. Number):** Maximum spend, pages, or runs allowed per shift: _____
[ ] **3. Expiry & Shadow Burn-In:** Badge dies at dawn; new automation passes clean shadow days first: _____
[ ] **4. Undo & Single-Writer Lock:** One writer commits at a time; rollback/revoke path tested awake: _____
[ ] **5. Named Plug-Puller (One Person):** Who can revoke access solo at 2 a.m. without asking a chat? _____

Card B — Connected Lock-Cluster Essays on Tygart Media

Thread / Rule	Live Tygart Media Essay
4-Min Counterfeit Audit	A Locked Door Keeps an Honest Man Honest
5-Line Overnight Card	The 2 A.M. Card & Bounded Approval Is the Hire
AuthN vs. AuthZ	A Key Is Not Permission & Who Holds the Keys
Draft-Only First Verb	We Put the Email on the Desk. Draft-Only Is the First Verb.
Many Readers, 1 Writer	Four Agents, One WordPress Lock & Chat Is a Room
Doors & Harness	Six AI Seats Governance & The Operator Is the Harness
Exhaust Series Siblings	Mad Libs · Two Droids · We Come From Water Damage

Verified Primary Authorities (All Checked Live Nov 20, 2026)

Why Prompts Alone Bend: Wallace et al., The Instruction Hierarchy: Training LLMs to Prioritize Privileged Instructions (arXiv:2404.13208, 2024); NIST AI 600-1 Generative AI Profile (2024, pp. 8, 13–15).

Keys, Scopes & Single-Writer Locks: Saltzer & Schroeder, The Protection of Information in Computer Systems (MIT/IEEE, 1975); NIST CSRC Authentication, Authorization & Least Privilege; IETF RFC 6749 OAuth 2.0 Scope; SQLite Write-Ahead Logging.

DAD'S LOCKED-DOOR RULE TO TAPE ABOVE YOUR DESK: Never tape a character essay to the steering wheel while leaving master keys on the van seat. Build the room with scoped verbs, hard caps, single-writer locks, and a human on every one-way door.