FREE EDITION · V1.0

The AI Adoption Triage

A 15-minute diagnostic for product teams shipping AI features.

12 failure modes. One symptom map. A 25-question self-audit that tells you which ones are eating your retention.

12 FAILURE MODES · 1 SYMPTOM MAP · 25-QUESTION SELF-AUDIT

The Symptom Map

Find the row that matches what you are seeing. If two rows match equally well, both diagnostics apply — adoption problems usually compound.

WHAT USERS DO

WHAT YOU ARE SEEING	FAILURE MODE
Signup → tried AI once → never came back	**01. First-Use Drop-off**
Users land on the prompt box, stare at it, type "help me" or leave	**02. Empty Prompt Paralysis**
"Wow, impressive!" — but they do not act on the output	**03. Trust Gap**
Generates results, but users heavily edit, reformat, or discard them	**04. Output Not Usable**
Output is plausible but wrong about user-specific facts	**05. Context Failure**
Users accept incorrect outputs in high-stakes workflows	**06. Overreliance Risk**
Users like suggestions but refuse agents, connectors, or automation	**07. Automation Anxiety**
Users keep correcting the same mistake and it never changes	**08. Correction Loop Breakdown**
Drop-off at upload, connector, or permission steps	**09. Privacy Hesitation**
First task succeeds, week-2 retention collapses	**10. Retention Without Habit**
Team cannot explain why AI quality improves or regresses	**11. Quality Ops Gap**
A few enthusiasts use it, but the team never standardizes	**12. Team Adoption Friction**

WHAT THE METRICS SHOW

METRIC PATTERN	MOST LIKELY FAILURE MODE
Signup → first-task funnel collapses before output	**01, 02, 09**
High generation count, low export / save / approve count	**03, 04**
High prompt rewrite rate	**02, 05**
Repeated identical corrections	**08**
High approval rate combined with low review time	**06**
Low connector or permission grant rate	**07, 09**
Day-7 or day-14 retention collapses	**10**
Surprise regressions after model or prompt changes	**11**
Pilot stalls before team rollout; admin questions go unanswered	**12**

WHAT USERS SAY

QUOTE YOU KEEP HEARING	FAILURE MODE
"I didn't know what to ask"	**02**
"It's impressive but I wouldn't send this without checking"	**03**

QUOTE YOU KEEP HEARING	FAILURE MODE
"I have to rewrite everything it gives me"	04
"You missed [obvious thing]"	05
"It approved it — I assumed it was right"	06
"I'm not comfortable letting it act on its own"	07
"I corrected this last week and it did it again"	08
"I don't know what happens to my data"	09
"I used it once, it was cool, I just haven't gone back"	10
"Our team can't agree on what good looks like"	11
"Our admin won't approve it for the team"	12

Use this page as your reference. Pin it, share it, route adoption conversations through it. The next two spreads describe each failure mode in detail.

The 12 Failure Modes

One paragraph each. This page names the problem. The full deck names the fix.

01. FIRST-USE DROP-OFF

Users sign up, run the AI once, and never return. The first session fails for a reason you have not yet labeled: unclear value, weak examples, too much setup, missing context, slow response, or an underwhelming first output. The dominant cause is rarely "users do not understand AI" — it is that the first concrete task, the first usable artifact, or the first proof-of-value moment is not concrete enough to anchor a second visit.

02. EMPTY PROMPT PARALYSIS

The product hands users a blank prompt box and asks them to invent the workflow. They stare at it, type "help me", and leave. The cost is not creativity — it is decision fatigue at the worst possible moment. Products that succeed here surface a small, recognizable set of starting tasks; products that fail keep the prompt blank and call it flexibility.

03. TRUST GAP

Users say the output is impressive. They do not act on it. They do not export it, send it, or stake a decision on it. The symptom looks like a quality problem but is usually a confidence problem: missing sources, no uncertainty signals, no second-glance affordances, no "show me why". High generation, low approval is the metric signature.

04. OUTPUT NOT USABLE

The AI produces something plausible but the user has to rebuild it before using it. Wrong format, wrong tone, wrong length, wrong structure, wrong place to paste it. The product is producing answers when the user needs artifacts — and the gap between the two becomes a tax that compounds every session.

05. CONTEXT FAILURE

The output is articulate, confident, and wrong about user-specific facts. The system has the wrong context, no context, stale context, or hidden retrieval the user cannot inspect. Users blame the AI; the diagnosis is almost always upstream — what the model could see, what it could not, and which silence it filled with confabulation.

06. OVERRELIANCE RISK

Users approve AI outputs faster than a human could read them. The dwell time on the review step is suspiciously short. The product makes "yes" easier than "check" and hides uncertainty behind confident prose. The cost lands later, in the consequences of an unreviewed decision — and lands on the user, not the product.

07. AUTOMATION ANXIETY

Users happily accept AI suggestions but refuse to let it act. Agents, connectors, integrations, automated workflows — all left switched off. The problem is rarely the capability and almost always the control surface: no reversibility, no visible scope, no preview of what is about to happen, no clear permission boundary.

08. CORRECTION LOOP BREAKDOWN

Users correct the AI. The product does not learn. The same mistake reappears next week, in the same shape, sometimes word-for-word. Feedback is captured at the wrong granularity, kept in the wrong layer, or never converted into a durable rule. Users notice within three corrections and stop correcting after five.

09. PRIVACY HESITATION

Drop-off at the upload step, the connector step, the "give us permission" step. The product is asking for data trust before it has earned it — sometimes before the user has seen anything work. The value-to-disclosure ratio is upside down: too much data asked for, too early, with too little clarity about what happens to it.

10. RETENTION WITHOUT HABIT

The first task succeeds. The user is delighted. They never come back. The output is a one-off — no saved artifact, no scheduled re-run, no surface inside an existing workflow, no reason to think about the product on a Tuesday morning. Retention does not collapse because the AI is bad; it collapses because the AI has no place to live.

11. QUALITY OPS GAP

Quality changes. The team cannot explain why. A model upgrade, a prompt tweak, a vendor change — and suddenly metrics move in directions nobody predicted. The product has no golden tasks, no rubric, no trace review, no regression dashboard. Improvements are vibes; regressions are firefights.

12. TEAM ADOPTION FRICTION

A few enthusiasts love it. The team never standardizes. The pilot stalls. Procurement asks questions nobody on the product side can answer: audit logs, policy controls, value reporting, admin scope. The product was built for an individual and is now being asked to behave like enterprise software — and the gap is visible at every step of the rollout.

The Self-Audit

Twenty-five yes/no questions. Roughly two per failure mode. Score honestly — every "no" is a category worth investigating.

> *How to score: answer Yes only if you can point to where this is true in the product today. "We have a ticket for it" is not yes. "It is on the roadmap" is not yes. If you cannot demo it on Monday, mark it No.*

FIRST-USE & ONBOARDING (01, 02)

- [ ] A new user can complete a meaningful AI task within their first session, without configuration.
- [ ] The first AI surface offers recognizable starting tasks — not just a blank prompt.
- [ ] We can show a new user a usable artifact before asking them to connect data.

TRUST & OUTPUT (03, 04)

- [ ] Every AI output shows its evidence, sources, or reasoning when the user asks.
- [ ] The output ships in a shape the user can use directly — not a draft they must rebuild.
- [ ] We track the gap between generated and acted on (export, approve, send, save).

CONTEXT & ACCURACY (05)

- [ ] Users can see and edit what context the AI is using before it answers.
- [ ] When the AI does not have the context to answer, it says so instead of guessing.

RISK & REVIEW (06)

- [ ] High-stakes outputs have a review gate the user cannot click past in under one second.
- [ ] The product surfaces uncertainty when confidence is low — not just confident prose at all times.

AGENTS & PERMISSIONS (07)

- [ ] Before any automated action, the user sees a preview of what is about to happen.
- [ ] Every AI-initiated action is reversible, or marked clearly as not.

CORRECTION & LEARNING (08)

- [ ] User corrections persist beyond the current session.
- [ ] A correction made by one team member is reflected when the next team member runs the same task.

DATA TRUST (09)

- [ ] The product asks for the minimum data needed to deliver the next visible value.
- [ ] Users can see, export, or delete what they have shared with the AI in under three clicks.

HABIT & RETENTION (10)

- [ ] At least one AI output type lives inside a recurring workflow (saved, scheduled, or triggered).
- [ ] Users have a reason to return to the AI surface other than starting over.

QUALITY OPERATIONS (11)

- [ ] We have a set of golden tasks we run before shipping any model, prompt, or context change.
- [ ] We can detect a quality regression in production within a week of it appearing.

TEAM ROLLOUT (12)

- [ ] An admin can see what AI is being used, by whom, with what data — on one page.
- [ ] We have a defensible answer to "what value did the AI deliver this quarter" for a buying team.

CROSS-CUTTING

- [ ] We have named the dominant adoption failure mode in this product within the last 90 days.
- [ ] Our product, design, and engineering teams would name the same failure mode if asked separately.
- [ ] We can point to a specific cohort whose behavior changed because of an AI-adoption design decision we shipped.

READING YOUR SCORE

- 22–25 yes — You are running a mature AI product. The remaining failure modes are worth a workshop, not a rewrite.
- 15–21 yes — Healthy foundation, two or three categories are leaking. Identify them and run targeted workshops.
- 8–14 yes — Several compounding failure modes. Start with the diagnostic at the top of the funnel (01, 02, 09) before touching the rest.
- 0–7 yes — The adoption problems are not subtle. Use the symptom map to choose one failure mode and fix it before adding features.

Where to go next

aiproduct.cards

You now have a name for the problem. The next question is what to do about it.

The AI Product Adoption Deck is the full working playbook this triage is drawn from:

- 12 Diagnostic cards — one per failure mode, with signals to inspect, questions to ask, and how to run the diagnostic with your team.
- 80 Action cards — tactical design patterns across ten stacks: product fit, first-use, trust, context, output quality, agents, correction loops, privacy, retention, team rollout. Each card names the pattern, the evidence, when to use it, when not to, and what to pair it with.
- 12 Workshop cards — facilitator-ready sessions that convert a diagnosis into shipped product decisions, copy, experiments, or specs.

Every card cross-references the others. Diagnose the failure mode, pull the action cards it points to, run the workshop that fits the problem class.

If this triage was useful, the deck is what it is a sample of.