

A FIELD GUIDE FOR F500 AI LEADERS

# The 5 *Fatal Flaws*

*That kill most AI strategies before they ship.*

# Why this guide *exists.*

---

Most F500 teams we talk to have run a year of AI pilots and still can't tell you which ones paid back. Not because the people are bad. Because nobody was paid to make the hard call.

This guide is the diagnostic NS Studio runs in paid engagements. The same five questions, the same kill criteria. Read it in 12 minutes. Apply it to your roadmap before next quarter's planning cycle.

If three or more of the flaws are live in your org right now, that's your next quarter's work. The cost of running a doomed pilot for another quarter isn't the budget. It's the credibility you spend defending it.

---

*"The most valuable thing an advisor can say is sometimes 'don't build that.'"*

# The five *flaws.*

- 
- 01 **The Vanity Pilot** \_\_\_\_\_  
A demo that cannot be defended in a one-page memo.
  - 02 **The Vendor Trap** \_\_\_\_\_  
Switching costs that compound faster than ROI.
  - 03 **The Strategy Theater** \_\_\_\_\_  
Decks without decisions.
  - 04 **The Shadow AI Tax** \_\_\_\_\_  
Six tools per department, zero accounting.
  - 05 **The Vitamin, Not Painkiller** \_\_\_\_\_  
Building things that aren't urgent enough to defend.

# 01 / The Vanity Pilot

---

## WHAT IT LOOKS LIKE

A polished demo. A live model in a sandbox. Screenshots circulated in a Slack channel called #ai-wins. Quarterly slides where the word "promising" appears more than twice. Nobody can name the number it moved.

## WHY IT'S FATAL

Pilots without defensible economics absorb the budget that would have funded the one pilot that worked. A vanity pilot left running for nine months at \$80K all-in is not a \$80K mistake — it is the four shippable bets you didn't take.

## DIAGNOSTIC

### THE VANITY PILOT TEST

`<role>`

Act as a CFO reviewing every active AI pilot in the company. You are paid to ask one question: which of these will I defend to the board, and which would I quietly kill if asked today.

`</role>`

`<task>`

For each pilot, write the one-page memo. Specific dollar value created or saved. Hostile question the board would ask. Survivable answer.

`</task>`

`<rules>`

- Pilot must produce a defensible dollar figure by week 6
- "Promising" is not a status – it's a stall
- If the memo cannot be written, the pilot is the answer

`</rules>`

`<output>`

Defendable list → Stall list → Kill list

`</output>`

## 02 / The Vendor Trap

---

### WHAT IT LOOKS LIKE

A renewal arrives. Procurement asks if the contract is worth it. Three teams say yes. Two say maybe. Nobody can describe what it would actually take to move off the vendor — data, prompts, evaluation rigs, internal champions, retraining. So the renewal goes through. Again.

### WHY IT'S FATAL

Switching costs compound. Every quarter on the wrong vendor adds integration debt, model-specific prompt scaffolding, and a small army of internal advocates whose careers are now tied to the choice. Eighteen months in, the migration is a six-figure project the org will never approve.

### DIAGNOSTIC

#### THE SWITCHING-COST INVENTORY

`<role>`

Act as a head of platform engineering who has migrated off two AI vendors before and knows exactly where the hidden costs live.

`</role>`

`<task>`

Run the 90-day exit drill for the top three AI vendors. Be specific. Surface every coupling that would slow a migration — data, evals, prompts, people.

`</task>`

`<rules>`

- Exit cost must be expressed in weeks and dollars
- Count internal champions as a switching cost
- If exit takes more than 90 days, lock-in is real

`</rules>`

`<output>`

Vendor → Exit weeks → Exit cost → Verdict

`</output>`

## 03 / The Strategy Theater

---

### WHAT IT LOOKS LIKE

A standing weekly. Working groups. A 47-slide deck with a maturity curve. Everyone agrees the slides are excellent. The deck is presented to a steering committee that approves the next deck. No engineer's calendar changes the following Monday.

### WHY IT'S FATAL

Strategy without decisions is a tax on the senior people who could be shipping. The real cost isn't the meetings — it's the false sense of motion that delays the moment someone has to say no to a flagship project that should never have started.

### DIAGNOSTIC

#### THE DECISION AUDIT

`<role>`

Act as a chief of staff with read access to every AI-related meeting on the leadership calendar last quarter. Your only job is to score them.

`</role>`

`<task>`

For each meeting, log the outcome: a decision (D), a deck (K), or a deferral (X). Compute the ratio. A healthy ratio is at least one D per three meetings.

`</task>`

`<rules>`

- A decision must be reversible only by another meeting
- A deck is not a decision, even when applauded
- Deferrals older than 30 days count as decisions not made

`</rules>`

`<output>`

D / K / X ratio → Worst-offender meetings → Cuts

`</output>`

# 04 / The Shadow AI Tax

---

## WHAT IT LOOKS LIKE

Marketing has four AI tools. Sales has three. Engineering has six. Procurement has logged none of them. Each is on a personal credit card, a pilot license, or a free-tier account. Total spend is unknown. Total data exposure is unknown. Nobody owns it.

## WHY IT'S FATAL

Shadow stacks are how mid-size enterprises end up with seven-figure AI bills nobody approved and a compliance posture nobody can describe. The cleanup is always more expensive than the early audit. The damage shows up the quarter a regulator asks.

## DIAGNOSTIC

### THE SHADOW STACK INVENTORY

`<role>`

Act as an internal auditor who is paid to find every AI tool currently in use, regardless of who approved it or which budget line it lives on.

`</role>`

`<task>`

Survey every team. Count tools. Count spend. Count owners. Map each tool to a data-sensitivity tier and a contract owner. Surface the unknowns.

`</task>`

`<rules>`

- Personal-card subscriptions count as company spend
- A tool with no owner is a tool to be cut
- If two teams use overlapping tools, only one survives

`</rules>`

`<output>`

Tool inventory → Spend → Owner → Cut list

`</output>`

# 05 / The Vitamin, Not Painkiller

## WHAT IT LOOKS LIKE

A project that everyone agrees is interesting. The deck describes it as "strategic." It has no internal customer who would escalate if it slipped. It has no external customer who has asked for it by name. It survives planning cycles by being adjacent to things that matter.

## WHY IT'S FATAL

Vitamin projects burn the most senior engineers in the org against problems no user feels. They survive because killing them feels rude. Six quarters in, the team is exhausted and the company has nothing to show but a maturity curve that moved one notch.

## DIAGNOSTIC

### THE URGENCY TEST

`<role>`

Act as the head of product who has to defend the roadmap to a board that will ask, for each project, who specifically is bleeding without it.

`</role>`

`<task>`

For every active AI initiative, name the person – internal or external – who would notice and complain by name if the project paused for six months.

`</task>`

`<rules>`

- The complainer must be nameable, not a persona
- "The whole company" is not a name
- No nameable complainer in 90 days = vitamin

`</rules>`

`<output>`

Project → Named complainer → Painkiller / Vitamin

`</output>`

# The kill *list.*

---

*"Every NS Studio engagement ends with a kill list. Not a deprioritize. Not a revisit-next-quarter. Kill."*

Explicit kills matter because the alternative is implicit drift. A project on the deprioritized list still consumes attention; it lingers in standups, in roadmaps, in the back-of-mind of the engineer who built it. A killed project frees the calendar.

The kill list is the most valuable artifact we produce because it's the one no internal politics will let you write yourself. Your VP can't kill the project a peer championed. Your peer can't kill the project the VP sponsored. We can.

## KILL LIST / Q4 - SAMPLE ENGAGEMENT

- × ~~Vendor X enterprise renewal — ROI undefended after 9 months~~
- × ~~Internal LLM fine-tuning project — duplicates a vendor offering~~
- × ~~AI center of excellence Slack channel — meetings without decisions~~
- × ~~Customer support copilot v2 — v1 hasn't shipped~~
- × ~~Generative search prototype — no nameable internal complainer~~

# Apply this to *your business.*

---

## 01 Score yourself.

For each of the five flaws, mark live, partly live, or not present in your org. Be honest. The diagnostic is worthless if scored generously.

## 02 Rank by cost.

Which flaw is bleeding the most money or credibility right now? Money is easier to count. Credibility is the one that matters more.

## 03 Run the prompt card.

Feed the prompt card to a senior team member — VP-level, not an analyst. Read the output literally. Do not soften it for the deck.

---

*"If three or more flaws scored 'live,' the work for next quarter is clear: stop, don't add."*

ENGAGEMENT

# When you want *a partner in this.*

We run this exact diagnostic for F500 teams as a paid engagement. Three engagements per quarter, founder-led, no junior consultants on client work.

ENGAGEMENT SHAPES

AI Audit 2-week sprint	\$4,500
Strategy Sprint 4 – 6 weeks	\$18,000 MOST POPULAR
Embedded Advisor 3-month minimum	\$9,500 / month

If the engagement isn't paying for itself by week two, we end it cleanly. The currency of an advisor is honesty.

*Book a 30-min call*

→ [nsstudiollc.com](https://nsstudiollc.com)

# Colophon.

---

## ABOUT THIS GUIDE

This is Field Guide 01, written by Simrat Bath, founder of NS Studio LLC. Updated quarterly. Distributed free, no email gate. The text is set in Instrument Serif and Inter; the prompt cards in JetBrains Mono.

## ABOUT NS STUDIO

NS Studio LLC is a boutique AI advisory based in California, working with Fortune 500 product and engineering teams. Certified Small Business 2045105. Three engagements per quarter, founder-led.

## CONTACT

[simratbath@gmail.com](mailto:simratbath@gmail.com)

[linkedin.com/in/simratbath](https://linkedin.com/in/simratbath)

[linkedin.com/company/ns-studio-llc](https://linkedin.com/company/ns-studio-llc)

[nsstudiollc.com](https://nsstudiollc.com)

---

<NS />

NS STUDIO LLC · CALIFORNIA · 2026