

SEMANTIC-FAILURE GRID

The Semantic-Failure Grid

Bring one AI workflow you ship. Find the hard cell: where syntax is green and meaning is wrong.

BRING one AI workflow whose output is released under a green status

1. Place each monitor on the grid

For each monitor or dashboard on your release path, name what it actually tests and which cell it can see. Most cannot see the hard cell - syntax green, meaning wrong.

Monitor or check	What it tests (format, latency, cost, policy, source, domain truth)	Cell it can see (syntax red / hard cell)	Can it block release?
Format / schema validator			
Latency or cost monitor			
Policy / toxicity filter			
Source or citation check			
Human sampling for meaning			
Other			

2. Name the relationship green is silently claiming to have checked

The verifier question is never generic. Pick the row for your domain and write the one relationship a passing output must actually satisfy. The prompts come from the six domains in the essay.

Domain	The relationship the verifier must resolve	Does any check verify it today? (Y / N)	Owner of the verifier
Legal - does the cited authority genuinely support the claim?			
Medical - does the response handle urgency and contraindications?			
Support - does the policy clause invoked match the customer's actual entitlement?			
Procurement - does the proposal satisfy each must-have requirement?			
Finance - does the transaction match the mandate and authorized threshold?			
Hiring - was the stated criterion applied consistently and lawfully across candidates?			

3. Close the hard cell

A green status can only honestly mean "no problem detected under the checks we chose to run". Confirm each statement is true of your workflow.

- We have written down the one relationship that, in our domain, green is silently claiming to have verified.

- A semantic verifier - deterministic check, off-family model, or human - actually resolves that relationship.

- The verifier is wired into the gate that releases the output, not run beside it.

- If a human performs the semantic check, that work is named, rostered and resourced - not an unbudgeted lunch-break habit.

- We can name the person or tool that checks the hard cell. If we cannot, we treat that as the gap.

4. The honest sentence

Write what your green status actually means, and who would catch a syntax-green, meaning-wrong output before it ships.

Companion worksheet to **Essay 16 · The Dashboard Is Green**, in the series **Architecting the AI Coworker**. · Dr Peter McCann Strain
· Fill this in against one real agent, action class or vendor. © 2026 Peter McCann Strain.

Series Companion + all 22 worksheets: **[Release_v12/Series_Companion.pdf](#)**