

AI Hallucinates Into Gaps.

Not randomly. Into the gaps you left.

A practitioner's field notes

Senthil Parameswaran · July 2026

THE CORE OBSERVATION

**AI doesn't
hallucinate
randomly.**

It hallucinates into
the gaps you left.

Missing requirements. Stale docs. Undesigned decisions. The agent fills them with something plausible — and it compiles.

Every project has two cost phases.

CHEAP PHASE

Upfront work

- ADR (architecture decision record)
- Requirement sentence
- Test vector
- Interface contract

Minutes to hours

EXPENSIVE PHASE

Rework

- Rework on wrong architecture
- Hallucinated API calls
- Confidently-wrong edits
- Debugging false-green tests

Days — and it repeats

1

Stale docs are worse than no docs.

Agents trust what they're given. An outdated spec isn't a minor inconvenience — it's a lie the agent acts on with full confidence. A live orientation script beats documentation every time.

2

A green test that skips the real path is the most dangerous failure mode.

AI writes tests consistent with the code it just wrote. If the code is wrong, the tests are wrong the same way — and they pass. At least one integration path must exercise the real system.

3

Architecture-first beats trial-and-error.

One week of AI trial-and-error produced working code that was architecturally wrong. An ADR with explicit tradeoffs and test vectors solved the same problem in a day. The difference was the clarity of the brief.

4

Two competing mental models produce one incoherent codebase.

When architecture exists in two conflicting forms — legacy code vs. new design — the agent mixes them. Make the canonical architecture explicit in one place before the agent starts.

5

Silent failures hide root causes.

AI agents generate optimistic code — errors get caught, logged, and swallowed. The system continues in a degraded state. When something finally fails visibly, the root cause is several steps back. Fail fast, loudly, with context.

6

Small verifiable truths drift constantly.

File paths, endpoints, token formats, schema shapes — the agent works from a snapshot. When those truths change, the agent doesn't know. A machine-readable source-of-truth catalogue, queried before every session, is the fix.

WHAT ACTUALLY WORKED

Live orientation

Query the real system before every session. Not docs — truth.

RCA before fix

Diagnose the assumption, not just the symptom. Fix the model first.

Architecture-first

Write the ADR before writing any code for complex components.

Real-device gate

Unit tests are necessary. They are not sufficient. Ship on the real target.

Machine-checked guards

Every production incident becomes a pre-push check. Incidents are expensive.

AI raises the stakes for engineering discipline.

It does not reduce them.

Disciplined team + AI: implementation cost collapses. Upfront investment stays constant. Net gain is large.

Undisciplined team + AI: wrong decisions are encoded faster, at greater scale, with greater consistency. Rework multiplies.

“
An AI agent is a brilliant contractor with no memory of yesterday and total faith in whatever you hand it.

Give it a clear blueprint: it builds fast.
Give it a lie: it builds that faithfully.

The brief is not overhead.

The brief is the work.

Full article at skaglobalpartners.com
Senthil Parameswaran · SKA Global Partners