

I Estimated Four Days. It Took Four Weeks.

Not because the AI was slow. Because nothing ever quite finished.

SWIPE →

Three traps behind the loop

01 Helpfulness churn

02 Vigilance fatigue

03 Lossy compression

80%

The point where a project shifts from creation to refinement.

AI is genuinely excellent at creation. Refinement is where the trouble starts — every session ends with a slightly different version of the same thing, never quite finished.

2

Revision attempts per section. Maximum.

If it still isn't right after two tries, the problem usually isn't the wording. It's the brief. Stop prompting — go fix the brief instead.

Revision 15

Round about here, I stop finding real problems and start inventing plausible-sounding ones just to justify another pass.



AI is a spectacular navigator, but it is a terrible captain.

THE CENTRAL ARGUMENT

Treat it as a specialist, not a captain

WRONG MODEL

Equal collaborator

Sets the direction

Iterates without limit

RIGHT MODEL

Fast, capable junior specialist

Works inside boundaries you set

Frozen brief, limited passes, defined handoff

Freeze the Architecture **Early.**

Lock scope, sections, and sequence before detailed content work begins. A global rewrite six revisions in is almost always a sign this step got skipped — not that the content is wrong.

When Do You Take the Wheel?

Once work is genuinely at 80–90%, stop iterating inside the AI interface and finish it by hand. That's exactly the terrain a human finishes fastest — and where letting AI keep going does the most damage.

Full piece:

terencekok.com/blog/ai-infinite-loop-why-4-day-project-took-4-weeks

The Four-Day Estimate Wasn't Wrong.

It was the estimate for the version of this project that had boundaries around it from the start.



Terence Kok · Enterprise AI Strategist

SAVE · COMMENT · FOLLOW