

An Agent Is Only As Reliable As The **Process It Replicates.**

Deploying agents without a validated process baseline doesn't accelerate productivity. It transfers accountability to a system that carries none.

SWIPE →

Four stages, in order

01 Build tool fluency

02 Identify the patterns

03 Specify the process

04 Build the agent against that specification

Guesswork

Dressed up as innovation — that's what agent-first looks like

Pattern recognition about where AI adds genuine value and where it introduces risk is the raw material a well-designed agent is built from. Skip it, and you're automating blind.

2-3

Years out is when today's tool-fluency work actually pays off

The professionals who derive the most sustained value from agents aren't rushing to deploy today. They're building the domain fluency that lets them govern agents reliably later.

What it actually looks like

- 01** Prompts that produce consistent, structured output
- 02** Knowing when a model's confidence is unreliable
- 03** Compressing low-value cognitive tasks
- 04** Understanding failure modes in your own context



Agents are not the destination. They are the industrialisation of a workflow that has already been validated by a capable, tool-literate human.

Terence Kok

Enterprise AI Strategist

Agent-first fails differently than fluency-first

AGENT-FIRST

No validated baseline

Errors are plausible-looking, hard to catch

Accountability transfers to a system with none

FLUENCY-FIRST

Process documented with precision


A human who can validate outputs

The agent has a testable baseline

Not obviously wrong. Plausibly wrong, at scale.

If the process is poorly understood or inconsistently executed, the agent fails in ways that are hard to detect. The outputs look right. They're applied at speed, across every instance, before anyone notices they weren't.

A structured pre-deployment check



The TRACE Agent Evaluation tool assesses readiness across five criteria:
terencekok.com/blog/stop-chasing-ai-agents-learn-tools-first

Augmentation before automation. Understanding before industrialisation.

The anxiety about falling behind is understandable. The rush to deploy agents without foundation is not justified.



Terence Kok · Enterprise AI Strategist

SAVE · COMMENT · FOLLOW