

Prove It. Don't Just Promise It.

AI is going into production faster than most organisations can build guardrails for it. Assurance is how you close that gap.

SWIPE →

Six practices working together

01 Performance checks, before and after deployment

02 Bias and fairness testing against real thresholds

03 Stress-testing against adversarial inputs and drift

04 Explainability for the decisions involved

05 Data lineage and provenance tracking

06 Human oversight where impact calls for it

ISO 42001

The AI management system standard that gives assurance its structure.

Organisations set up, run, and keep improving how they govern AI in their own context — paired with ISO/IEC 42005 for lifecycle impact assessment.

Probabilistic.

Unlike normal software, AI output can shift even when the code doesn't.

A system that passed every check at launch can quietly drift off course in production. Continuous monitoring is what catches that — not a one-off sign-off.

A Black Box.

Foundation models, third-party fine-tuning, and AI features buried inside procured software mean a lot of your risk exposure comes from outside your own engineering team.



Assurance is how you prove an AI system does what it's supposed to do, with evidence, not good intentions.

THE CENTRAL ARGUMENT

The timeline is compressing either way

RETROFIT LATER

Costs far more, more disruption

Fixed on the regulator's clock

A box-tick beside the "real" work

BUILT IN FROM THE START

Risk-ranked inventory of every system

Continuous monitoring, real numbers

An operating function, not a checkbox

The Cost of Getting It Wrong Isn't Symmetrical.

One serious failure — a biased decision, a hallucinated output where it really shouldn't happen, a compromised model — can do damage completely out of proportion to what it would have cost to catch it early.

Could You Prove It to a Regulator Today?

Get evidence in place now, and you've got a solid baseline to stand on when scrutiny comes. Leave it too late, and you'll build it under pressure, after something has already gone wrong.

Full piece:

terencekok.com/blog/ai-assurance-why-you-need-it-now-not-later

Evidence. Not Good Intentions.

That's the whole test. Cover your highest-risk AI systems now, before the regulator's clock starts running instead of yours.



Terence Kok · Enterprise AI Strategist

SAVE · COMMENT · FOLLOW