

Not Every Process Is Ready for an Agent. **Five Traits Tell You Which Are.**

40% of enterprise apps will carry task-specific agents by end of 2026.
The question is which of your processes are structurally ready.

SWIPE →

What this deck covers

- 01** What distinguishes agentic AI from prior automation
 - 02** Five structural traits of agent-ready processes
 - 03** Where agentic AI is already delivering, by domain
 - 04** Governance: Singapore's four-principle framework
-

40%

Of enterprise apps will carry agents by 2026

Up from under 5% in 2025. McKinsey estimates \$450–650B in additional annual revenue by 2030 across advanced industries alone.

Source: Gartner; McKinsey

90%

Of SOC's are overwhelmed by alert backlogs

The domain where agentic AI delivers the most acute value. Autonomous triage and remediation have shown up to 90% reductions in response times.

Source: industry SOC survey data

307%

Three-year ROI from agentic AI in financial operations, in a Forrester Total Economic Impact study — plus US\$3.4 million in incremental revenue.

Source: Forrester Total Economic Impact study



Traditional automation breaks when conditions deviate from scripted rules; generative AI produces outputs but doesn't act on them; agentic AI reasons, acts, monitors, and self-corrects.

Terence Kok

Enterprise AI Strategist

The screening framework, at a glance

NOT AGENT-READY

Ad hoc, low-frequency, judgement-only

Single system, no integration needed

No measurable outcome

AGENT-READY

Repetitive, rules-governed, high-frequency

Multi-step, cross-system

Measurable, tolerates adaptive execution

The constraint isn't capability. **It's readiness.**

The technology is production-ready. What's missing is institutional readiness — the willingness to redesign process architectures, establish adaptive governance, and commit to continuous measurement against quantified baselines, not enthusiasm.

Autonomy scales with risk, not ambition.

LOW-RISK	Fully autonomous execution, continuously monitored.
MEDIUM-RISK	Human-on-the-loop — oversight without blocking every step.
HIGH-RISK	Hard human-in-the-loop gate before any irreversible action.

Singapore's Model AI Governance Framework for Agentic AI:
terencekok.com/blog/critical-tasks-suitable-agentic-ai

Score your next use case **before you build it.**

Five structural traits separate agent-ready processes from expensive experiments.



Terence Kok · Enterprise AI Strategist

SAVE · COMMENT · FOLLOW