Terence Kok · Framework reference brief · Version 1.1 · First published 15 June 2026 · Revised 12 September 2026 · CC BY-NC 4.0 · terencekok.com/frameworks/governance/

AI GOVERNANCE

Four-Question AI Governance Baseline

Derived from the governance framework applied across Singapore IMDA, Saudi NDMO, and Oman TRA regulatory environments at Meinhardt Group. Referenced in the Group Annual Sustainability Report.

Most organisations treating AI governance as a documentation exercise discover the gap when something goes wrong. The four questions below are diagnostic, not procedural: they expose structural governance gaps before deployment, not after incident. The complexity of the deployment is irrelevant. The questions are the same whether you are deploying across three regulatory jurisdictions or adding a single AI assistant to one team.

REFERENCE DIAGRAM

Four diagnostic questions around a single deployment. Each establishes one structural control; any question left unanswered means the deployment runs on trust.

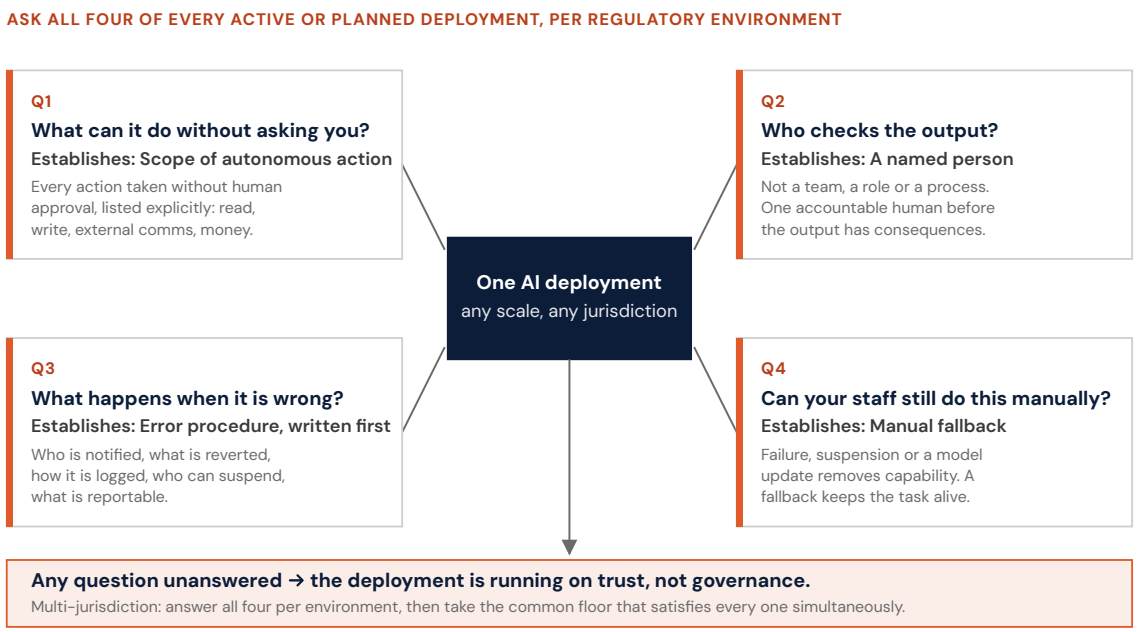

Q1

What can it do without asking you?

Maps the scope of autonomous action. Every action the AI system can take without human approval must be explicitly defined, not assumed to be limited. Gaps in this list are gaps in your governance. For agentic systems, this includes read access, write access, external communications, and financial commitments.

Q2

Who checks the output?

Names a human: not a team, a role, or a process, but a named person accountable for reviewing AI output before it produces consequences.

"The system checks itself" is not an answer. If nobody is named, governance does not exist regardless of what the policy document says.

Q3

What happens when it is wrong?

Documents the error response procedure before the first error occurs. Covers: who is notified, what is reverted or corrected, how the incident is logged, who has authority to suspend the system, and what constitutes a reportable event under applicable regulatory frameworks. Absence of a documented procedure means the first error will be handled inconsistently.

Q4

Can your staff still do this manually?

Tests operational resilience. AI system failures, regulatory suspensions, or model updates can remove capability without warning. If staff cannot perform the task manually, the organisation has created a critical dependency without a fallback. This is a governance risk in regulated environments and an operational risk in any environment.

Application: Run these four questions against every active or planned AI deployment in your organisation. Any deployment that cannot answer all four is not governed. It is running on trust. For multi-jurisdiction deployments, apply the questions against each regulatory environment separately, then identify the common governance floor that satisfies all of them simultaneously.

CITE THIS FRAMEWORK

VERSION	FIRST PUBLISHED	LAST REVISED	LICENCE
1.1	15 June 2026	12 September 2026	CC BY-NC 4.0

DOI
10.5281/zenodo.22722219

APA | Harvard | Chicago | BibTeX

Kok, T. (2026). Four-Question AI Governance Baseline (version 1.1) [Framework reference]. terencekok.com. https://doi.org/10.5281/zenodo.22722219

Copy citation