

A regulator drew columns and called them pillars.

IMDA published the first **national** framework written specifically for agentic systems. It organises technical controls by lifecycle stage — design, development, pre-deployment, deployment. Those are layers. But two of its four pillars are not stages at all. They run across every stage at once. **A regulator built the columns shape without naming it.**

MGF PILLARWHAT IT REQUIRESLAYER OR COLUMN

Assess and bound the risks
pillar 1Likelihood × impact, judged before deployment. v1.5 adds systemic and multi-agent risk: agent sprawl, collaborative failure, emergent behaviour, third-party-agent visibility.**LAYER**

Make humans meaningfully accountable
pillar 2Responsibility allocated inside and outside the org. **Mandatory human approval for irreversible, high-stakes or atypical actions.** Oversight complemented by real-time monitoring.**COLUMN**

Implement technical controls
pillar 3Prefer **structural, rule-based** controls — block the tool at system level — over prompt-layer instruction not to use it. Runtime controls layered on top.**LAYER**

Enable end-user responsibility
pillar 4Inform users of agent capabilities and limits. Train them to oversee agents. Tailor the information to different user needs.**COLUMN**

WHY THIS IS THE EXTERNAL PROOF POINT, NOT A FOOTNOTE
When a director pushes back that human oversight is *a deployment-stage gate*, the answer is no longer an argument. **A national AI regulator structured its own framework to put human accountability across all four lifecycle stages rather than at one of them** — and did it independently. Voluntary and non-binding, and organisations stay legally accountable for their agents’ actions whether they adopt it or not.

WHAT IT COVERS THAT THE BINDING REGIMES HAVE NOT CAUGHT UP TO

- 01
Agent sprawl — how many are running, and who counted.
- 02
Collaborative failure — miscoordination, conflict, collusion between agents.
- 03
Structural over prompt-layer control — block the tool, do not ask the agent nicely.

Every framework that takes agents seriously ends up separating the stages from the properties.

Source: Model AI Governance Framework for Agentic AI, Infocomm Media Development Authority, Singapore. Launched 22 January 2026 at Davos; **v1.5 published 20 May 2026** incorporating feedback from 60+ organisations. Builds on IMDA’s 2019 Model AI Governance Framework and the 2024 generative-AI edition. **Voluntary and non-binding.** Described here as the first national framework written specifically for agentic AI — IMDA’s own characterisation — not as an absolute first. Layer/column assignments are this sheet’s reading, not IMDA’s framing.

FOLLOW FREDRIK LINDSTROM FOR MORE

Fredrik Lindstrom

AI and cybersecurity governance for boards, CISOs and senior leaders. 25+ years in enterprise security.

A practitioner to guide you.