



Center of Excellence (CoE) Preliminary Charter

Securing Singapore's AI Future · Enabling Confident, Governed AI Deployment

Grounded in Proportionality · Practicality · Independence

Managing the Delicate Governance Balance

UNDER-GOVERNED CREATES RISK

AI Deployed Without Guardrails

Enterprises adopt AI in ways that expose customers, staff, and society to avoidable harm — without adequate technical controls or accountability.

VS

OVER-GOVERNANCE CURTAILS INNOVATION

Innovation Paralysed by Compliance

Blanket compliance processes designed for high-stakes scenarios are applied to low-risk tools — creating friction that stalls innovation and drives AI initiatives underground.

SafeAI.sg is established to close this gap — Singapore's industry-led Centre of Excellence for AI technical safety and practical governance, enabling confident AI deployment at the appropriate level of rigour.

Mission & Five Delivery Pillars

The core mission of SafeAI.sg is to make AI deployment safer, more confident, and more economically rational — providing a risk-weighted governance framework, curated technical controls, and a trusted practitioner community that enables Singapore enterprises, agencies, and technology teams to deploy AI with confidence at the right level of rigour.

01

Use-Case Centric Risk Assessment Toolkit

Self-service scoring toolkit, freely accessible online. Any organisation determines the right governance tier for any AI use case in under 30 minutes.

02

Technology Associate & Showcase Programme

Technology companies demonstrate how AI governance controls work in real production environments. SafeAI.sg does not endorse, certify, or validate any Associate technology.

03

Professional Certification & Skills Development

Co-developed with Singapore's institutes of higher learning. Four certification levels from career conversion (Level 0) through to Chief Architect (Level 3).

04

Agentic Runtime Observability & Governance

Govern declared agent intent before action is taken — not after. Addressing three critical gaps in pre-execution interception, trajectory monitoring, and context integrity.

05

AI FinOps & Cost Governance

Token economics as a governance concern — from budget controls and context hygiene to denial-of-wallet attack detection and AI cost anomaly as a security signal.

Operating Principles

Proportionality

Governance commensurate with risk. No AI use case bears more governance friction than its actual exposure warrants.

Practicality

Every output designed for operational implementation by a working technology team — not merely for policy documentation.

Independence

SafeAI.sg does not regulate, certify, or endorse. Credibility rests on quality of analysis and trust of community.

Inclusivity

Frameworks accessible to mid-market enterprises, public agencies, and technology teams of all sizes.

Continuous Improvement

Committed to a continuous update cycle for UARS, control recommendations, and the tool reference library.

Transparency

Methodology, framework development process, and risk band calibrations publicly documented and open to challenge.

CHARTER ADMINISTRATION

Operational Foundation for Confident AI Deployment

Annual Charter Review

Reviewed by the Advisory Council at minimum annually and updated to reflect material changes in scope, structure, or operating environment.

Amendment Process

Proposed by executive team or any Council member. Circulated 30 days before formal meeting; requires majority agreement to take effect.

Effective Date

Takes effect upon adoption by the founding Advisory Council and remains in force until superseded by a subsequently approved version.

CONTACT US


Contribute to SafeAI.sg

**Laura Flores**

Assistant Chief Secretariat
laura@wahspark.com

**David Chin**

Chief Secretariat
David.chin@wahspark.com

A large, semi-transparent teal circle is positioned on the right side of the slide, partially overlapping the dark background.