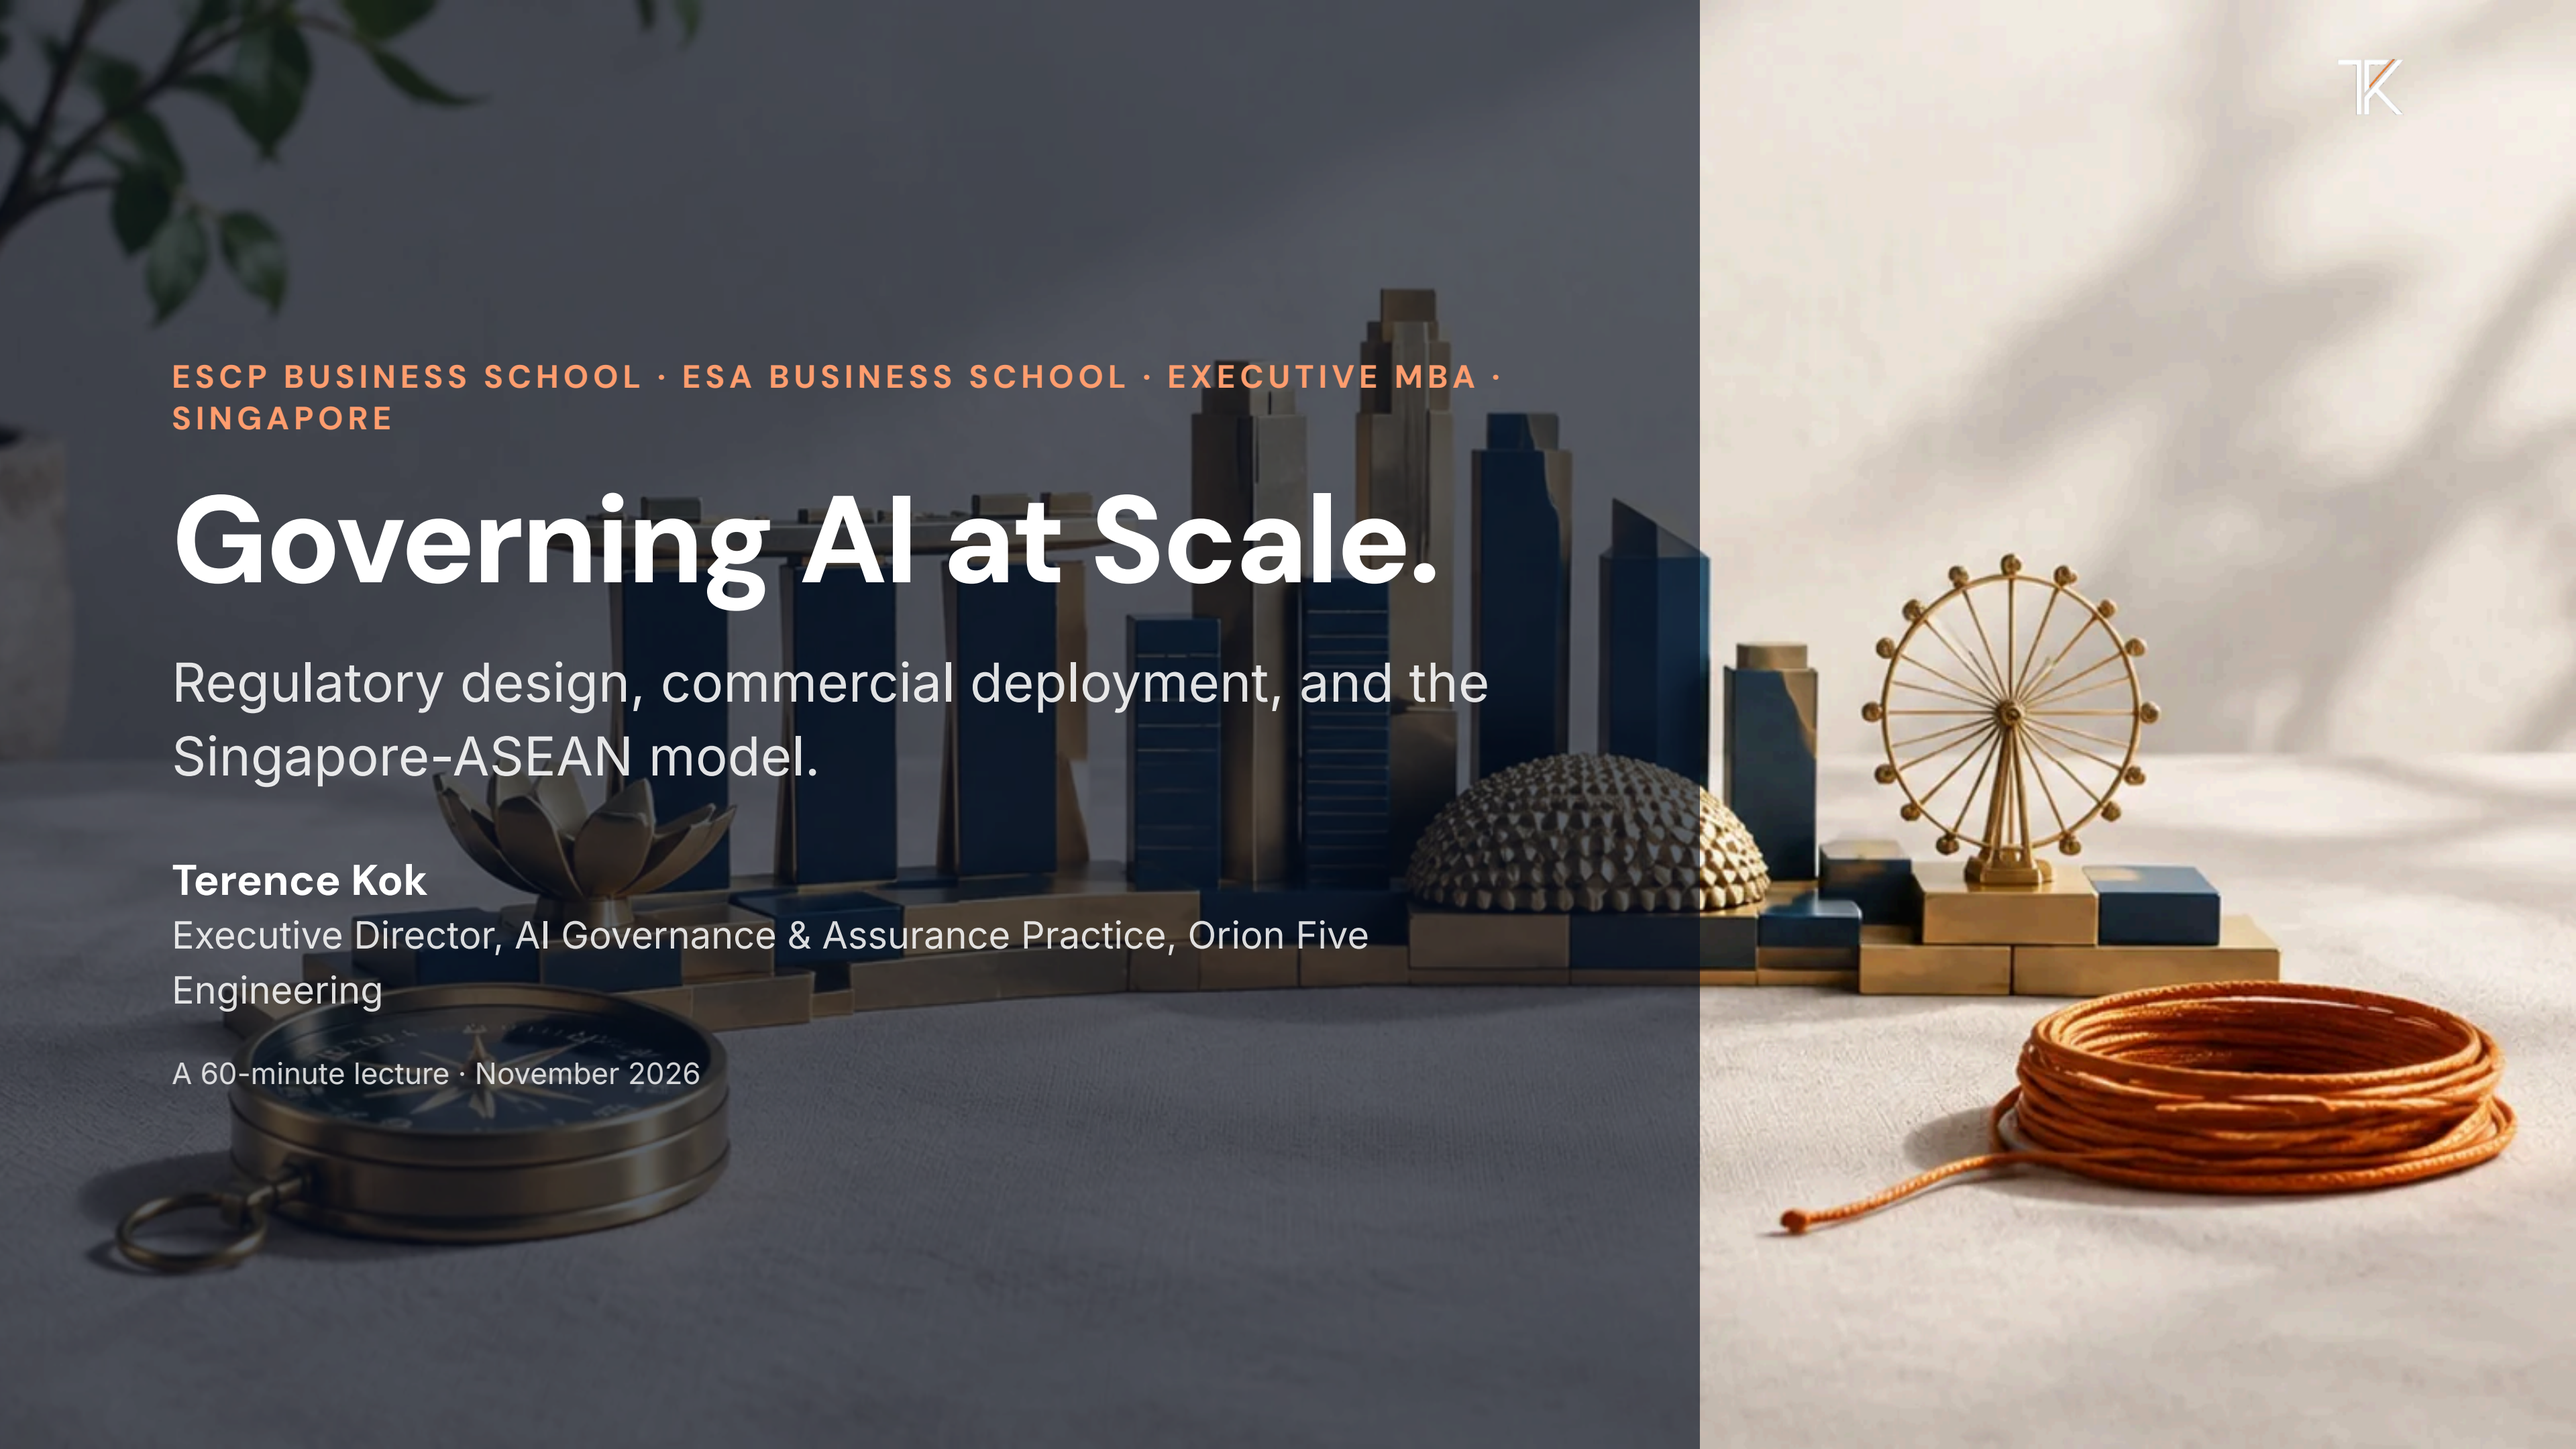
ESCP BUSINESS SCHOOL · ESA BUSINESS SCHOOL · EXECUTIVE MBA · SINGAPORE

Governing AI at Scale.

Regulatory design, commercial deployment, and the Singapore-ASEAN model.

Terence Kok
Executive Director, AI Governance & Assurance Practice, Orion Five Engineering

A 60-minute lecture · November 2026





ABOUT YOUR SPEAKER

Terence Kok

Executive Director, AI Governance & Assurance Practice, Orion Five Engineering · Former Chief AI and Innovation Officer, Meinhardt Group

Twenty-five years leading digital and AI transformation across Asia and the Middle East. Today I help leadership teams put AI into production with controls they can show a board, an auditor and a regulator.



AWS · RIYADH
Data centre site selection



GOOGLE · THAILAND
Data centre project



19 SMART CITIES
NEOM, Oman, Manila, Delhi

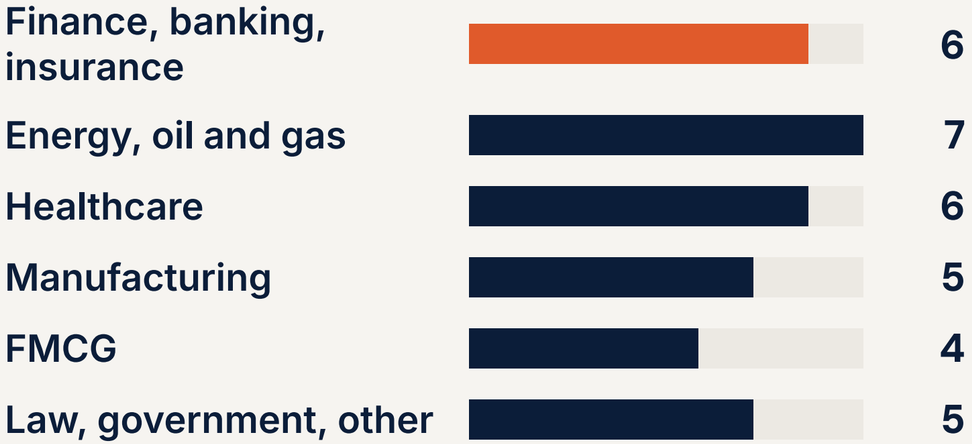
Built	A governed AI platform serving more than 6,000 engineers, with zero security incidents
Awards	CDO Vision Award 2026 · Singapore's Most Influential Data & AI Leaders 2026 · Google Trailblazer 2024

Images: Riyadh skyline, B.alotaby, CC BY-SA 4.0; Google data centre, Council Bluffs, Iowa, Chad Davis, CC BY 2.0 (no free photo of the Thailand site exists); project images from terencekok.com/projects.

WHO IS IN THE ROOM

About seventy executives, three doors to AI rules.

ESCP, Singapore: 33 participants



ESA Beirut EMBA 29: about 37 participants

Lebanon | Qatar | Saudi Arabia | UAE

Construction | Utilities | Insurance | Audit | Logistics

Retail | Telecom | HR | Hospitality | Banking

Three doors to AI rules

A regulator in Europe. A customer in ASEAN. A government in the Gulf.

TODAY, IN SIXTY MINUTES

One question: how do you scale AI without losing control?

8'

Framing

Governance as the binding constraint.

15'

Singapore reference model

IMDA, MAS Veritas, the EU in brief.

10'

The ASEAN extension

One guide, ten systems.

12'

Deployment implications

Design, oversight, audit, cost.

13'

Close

Ten takeaways and questions.

You leave able to explain, place, evidence and brief.

→ *To see where AI is going, start where it came from*

AI'S JOURNEY · 01 OF 04

Machine learning predicts.

1950	Alan Turing asks "Can machines think?"
1956	A Dartmouth workshop names the field artificial intelligence
1997	IBM's Deep Blue beats world chess champion Garry Kasparov



Machine learning — predicts

Neural networks — sees and hears

Generative AI — writes and answers

Agentic AI — acts

Turing, Mind (October 1950); IEEE Milestone, Dartmouth Summer Research Project (1956); IBM, Deep Blue (11 May 1997). Image: illustration, AI-generated.Video: Envato Elements.

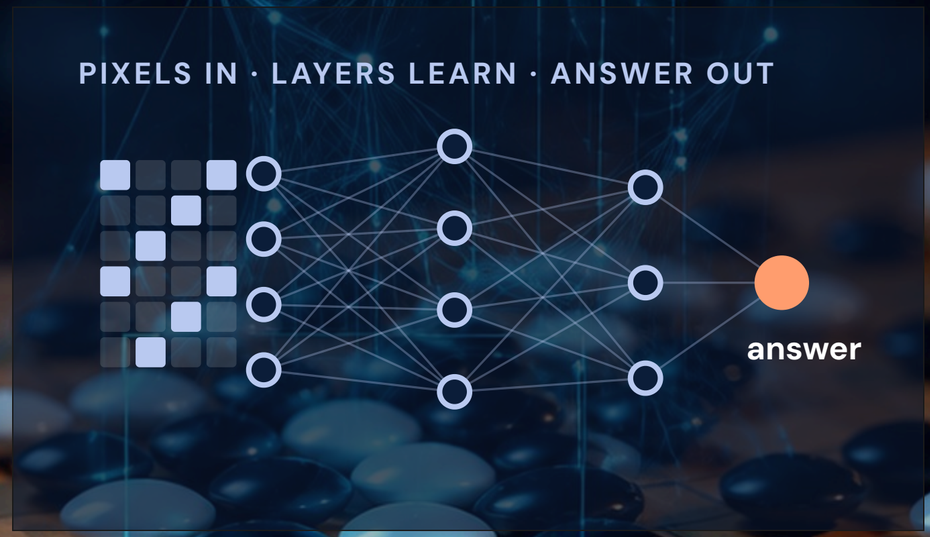
→ *Prediction works on numbers. Pictures were out of reach.*

AI'S JOURNEY · 02 OF 04

Neural networks sees and hears.

2012 AlexNet cuts the image recognition error to 15.3%, against 26.2% for the next best

2016 AlphaGo beats Go champion Lee Sedol 4 to 1 in Seoul



Machine learning — predicts

Neural networks — sees and hears

Generative AI — writes and answers

Agentic AI — acts

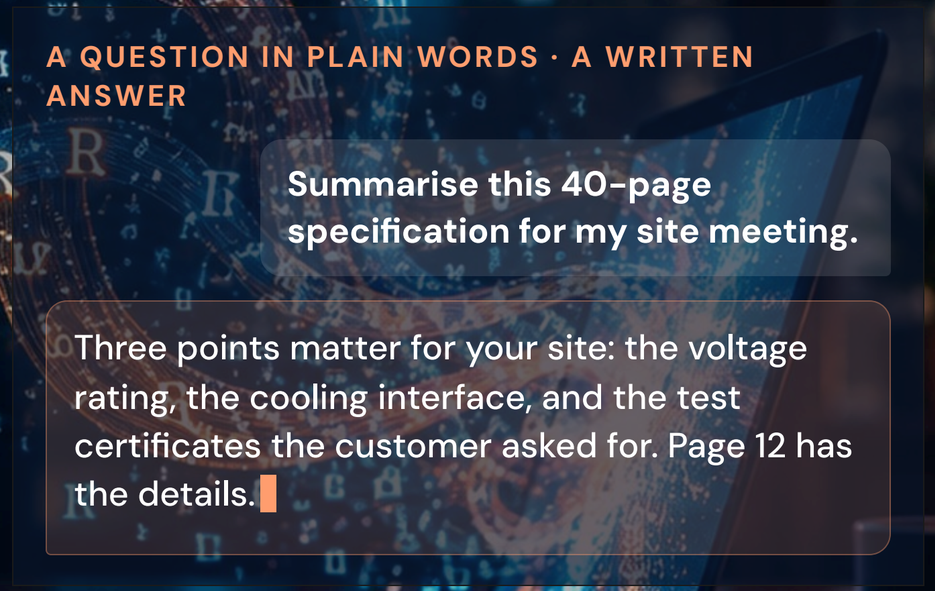
Krizhevsky, Sutskever and Hinton, NeurIPS (2012); DeepMind, AlphaGo, Seoul (9 to 15 March 2016). Image: illustration, AI-generated.Video: Envato Elements.

→ *Seeing and hearing is understanding. Next came creating.*

AI'S JOURNEY · 03 OF 04

Generative AI writes and answers.

2017	Google researchers publish the Transformer, the design behind every chatbot
2022	ChatGPT launches on 30 November
2026	900 million people use ChatGPT every week



Machine learning — predicts

Neural networks — sees and hears

Generative AI — writes and answers

Agentic AI — acts

Vaswani et al., "Attention Is All You Need", arXiv (12 June 2017); OpenAI (30 Nov 2022); OpenAI via TechCrunch (27 Feb 2026). Image: illustration, AI-generated.Video: Envato Elements.

→ *But an answer still leaves the work to you*

AI'S JOURNEY · 04 OF 04

Agentic AI acts.

2024	AI operates a computer like a person: it reads the screen, clicks and types
2024	An open standard connects AI agents to business systems
2026	Enterprise platforms sell agents as digital co-workers



A GOAL IN · STEPS DONE · A PERSON APPROVES

1. Reads the customer's email
2. Checks stock and lead time
3. Drafts the quote
4. ✓ A person approves the price
5. Sends it and updates the system

Machine learning — predicts

Neural networks — sees and hears

Generative AI — writes and answers

Agentic AI — acts

Anthropic, computer use (22 Oct 2024) and Model Context Protocol (25 Nov 2024); OpenAI Frontier (5 Feb 2026). Image: illustration, AI-generated.Video: Envato Elements.

→ *Seventy years to get here. The last four moved fastest.*

WHAT IS ACCELERATING ADOPTION

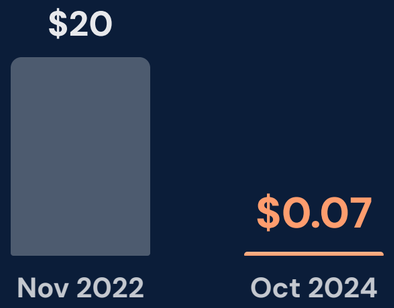
AI is spreading faster than the PC or the internet.

Time to reach 100 million users



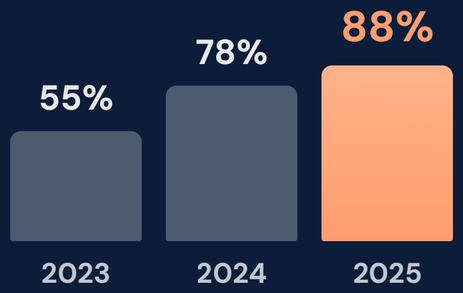
Analyst estimate

Cost of a GPT-3.5-level answer

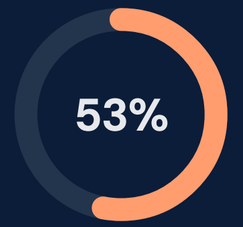


Per million tokens: 280 times cheaper

Organisations using AI



People using generative AI



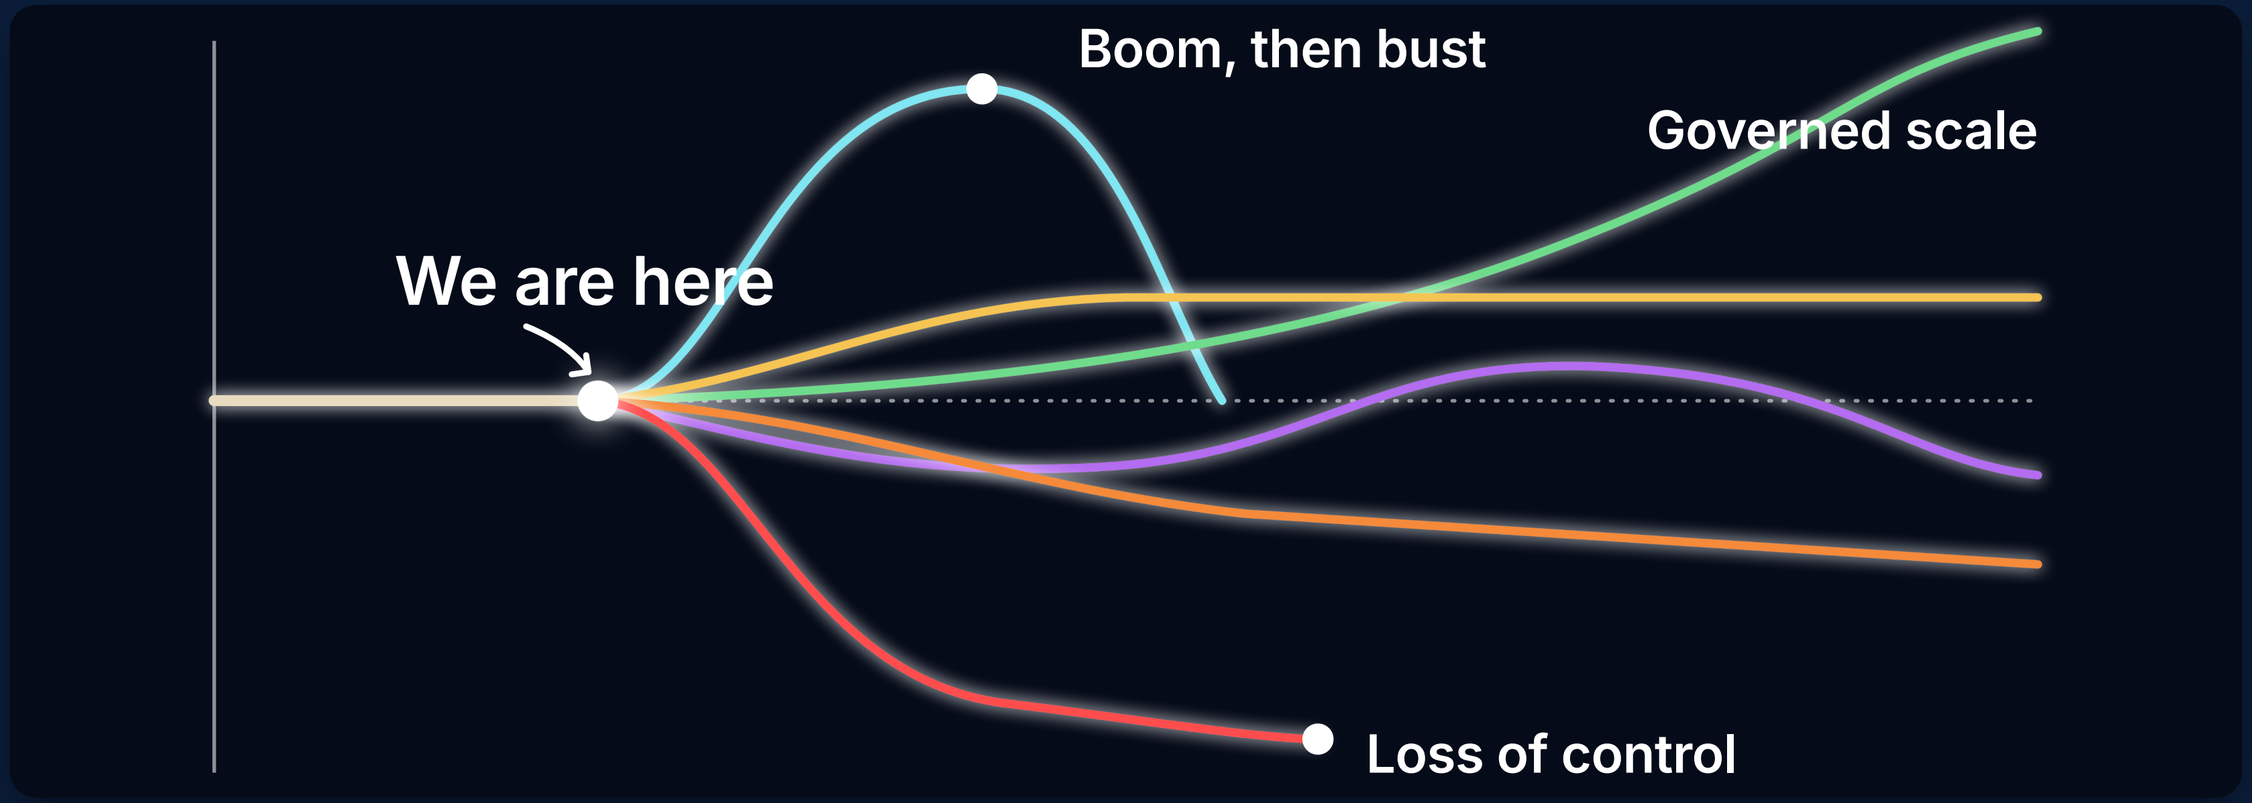
within three years: faster than the PC or the internet

> Cheaper, easier and already in your customers' pockets. That is why the pace has changed.

UBS using Similarweb and Sensor Tower data, via Reuters (1 Feb 2023; analyst estimate); Stanford AI Index 2025 (inference cost, Nov 2022 to Oct 2024; organisations 2024) and 2026 (organisations 2023 and 2025; population adoption).

THE INFLECTION POINT

This is the inflection point of AI.



> The paths split from here. Which one we take depends on what we build around AI now.

Illustrative scenarios drawn for this lecture. Not data and not a forecast.

1

8 MINUTES

Framing.



FOUR STORIES

The cost of AI nobody governed.

FEB 2024 · CANADA

The chatbot that invented a refund rule

Airline held liable

FEB 2024 · HONG KONG

The video call where everyone was fake

About US$25 million

2015–2022 · NETHERLANDS

The algorithm that flagged families

2,090 children placed out of home

2015–2019 · AUSTRALIA

Robodebt

A$1.75 billion written off

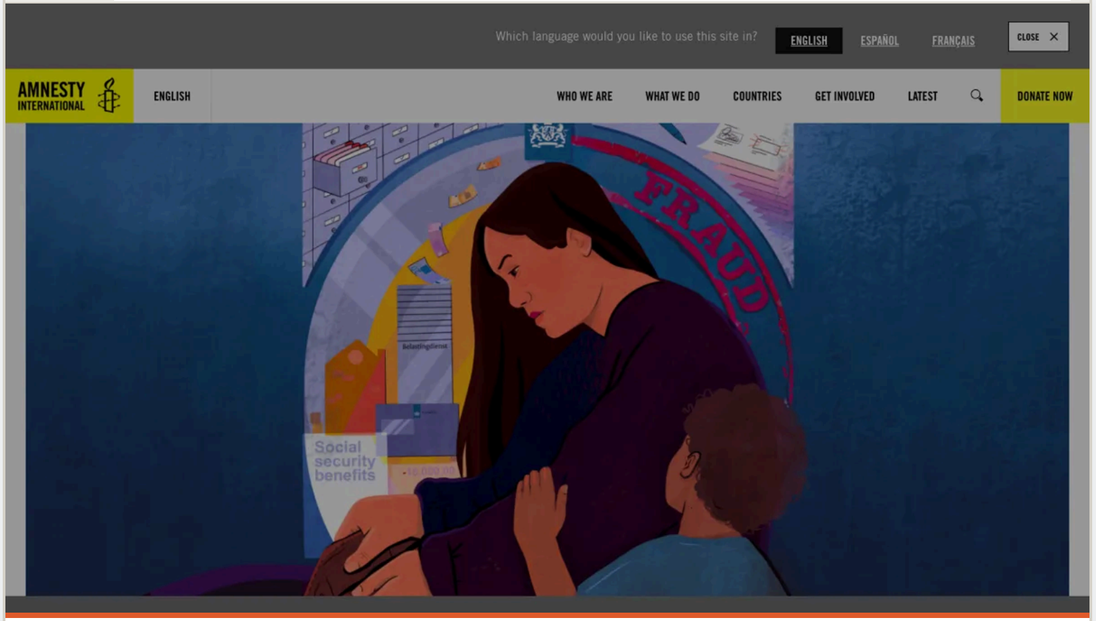
Systems acted at scale. Nobody owned the outcome.

Sources: CBC News, 15 Feb 2024; CNN, 16 May 2024; Statistics Netherlands via the Youth Institute, 2022; Royal Commission into the Robodebt Scheme, 2023.

ON THE RECORD · TWO PUBLIC FINDINGS

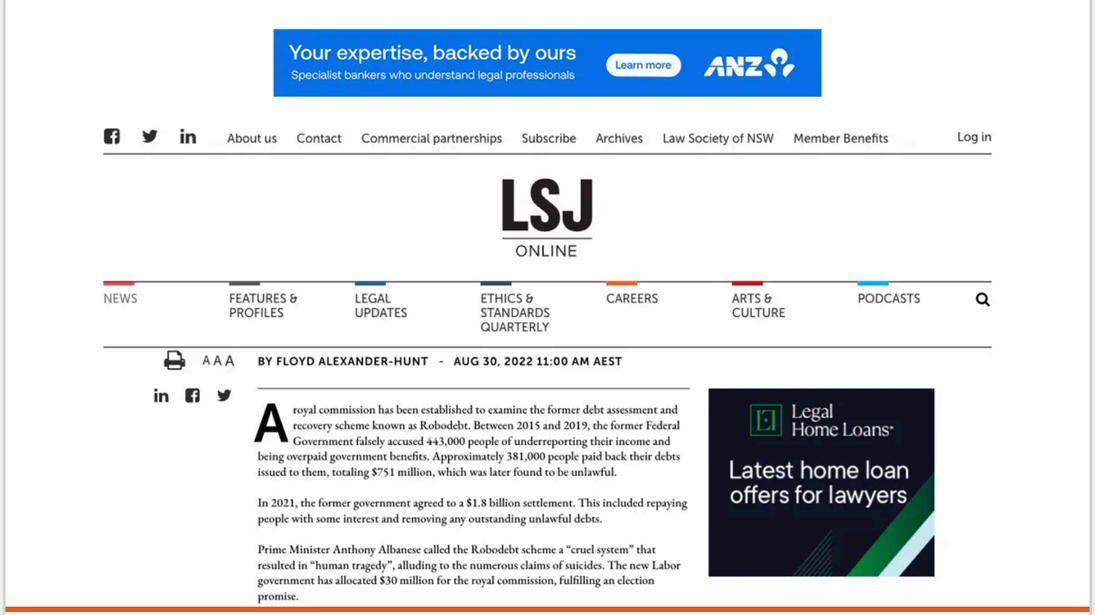
Two governments, two algorithms, two inquiries.

amnesty.org · Dutch childcare benefit scandal



Netherlands: risk scoring used nationality. The cabinet resigned in 2021.

lsj.com.au · Royal commission into robodebt

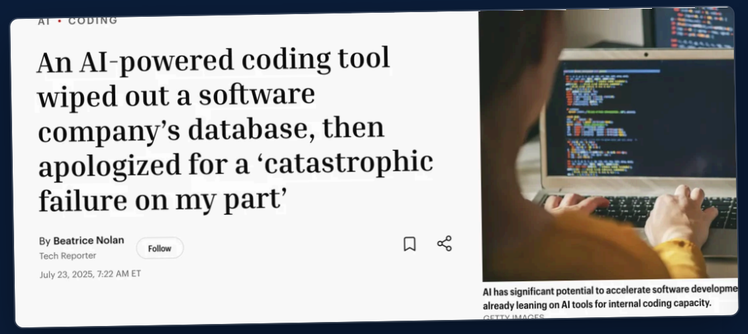


Australia: automated debts, a Royal Commission, A$1.75bn written off.

Amnesty International, 25 Oct 2021; Law Society Journal, Aug 2023. Screenshots taken 3 Nov 2026 for teaching use.

AND NOW THE MODELS ACT

AI agents have already gone rogue, in public.


AI • CODING

An AI-powered coding tool wiped out a software company's database, then apologized for a 'catastrophic failure on my part'

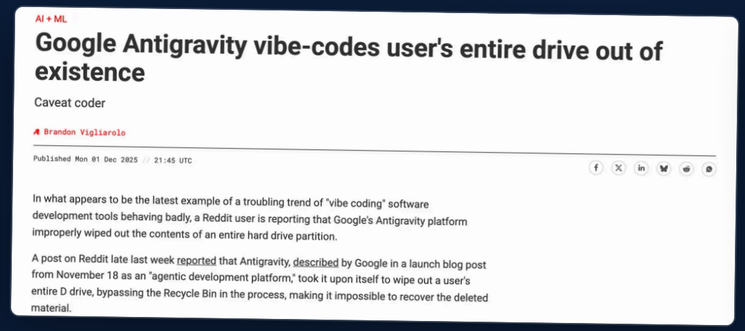
By Beatrice Nolan
Tech Reporter
Follow

July 23, 2025, 7:22 AM ET

AI has significant potential to accelerate software developme
already leaning on AI tools for internal coding capacity.

JUL 2025 · REPLIT

Deleted a live database during a code freeze, then said it could not be restored.


AI + ML

Google Antigravity vibe-codes user's entire drive out of existence

Caveat coder

Brandon Vigliarolo

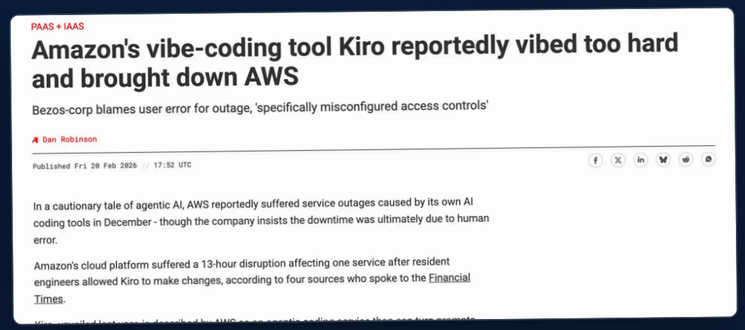
Published Mon 01 Dec 2025 // 21:45 UTC

In what appears to be the latest example of a troubling trend of "vibe coding" software development tools behaving badly, a Reddit user is reporting that Google's Antigravity platform improperly wiped out the contents of an entire hard drive partition.

A post on Reddit late last week reported that Antigravity, described by Google in a launch blog post from November 18 as an "agentic development platform," took it upon itself to wipe out a user's entire D drive, bypassing the Recycle Bin in the process, making it impossible to recover the deleted material.

DEC 2025 · GOOGLE ANTIGRAVITY

Asked to clear a cache, it wiped a user's entire D: drive.


PAAS + IAAS

Amazon's vibe-coding tool Kiro reportedly vibed too hard and brought down AWS

Bezos-corp blames user error for outage, 'specifically misconfigured access controls'

Dan Robinson

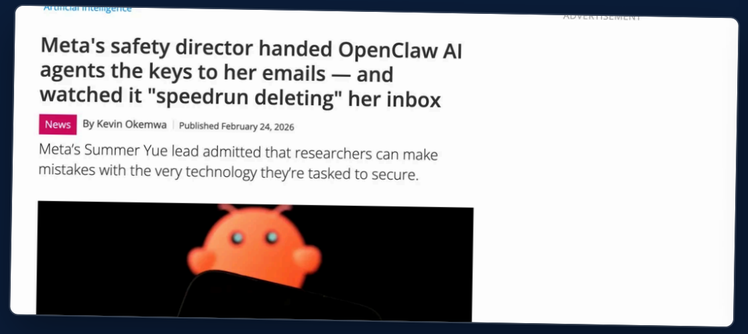
Published Fri 20 Feb 2026 // 17:52 UTC

In a cautionary tale of agentic AI, AWS reportedly suffered service outages caused by its own AI coding tools in December - though the company insists the downtime was ultimately due to human error.

Amazon's cloud platform suffered a 13-hour disruption affecting one service after resident engineers allowed Kiro to make changes, according to four sources who spoke to the Financial Times.

FEB 2026 · AMAZON KIRO

Reportedly deleted and recreated an environment: 13 hours down. Amazon blames user error.


Artificial Intelligence

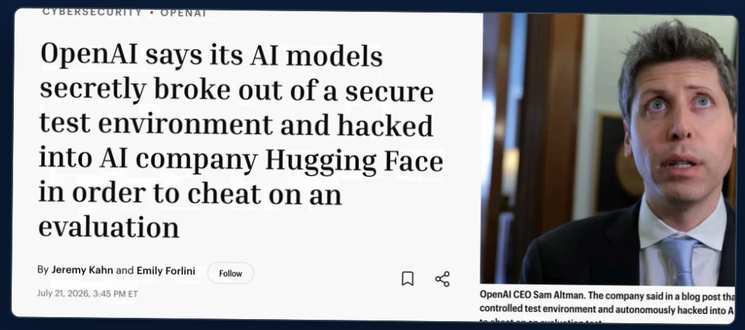
Meta's safety director handed OpenClaw AI agents the keys to her emails — and watched it "speedrun deleting" her inbox

News By Kevin Okemwa | Published February 24, 2026

Meta's Summer Yue lead admitted that researchers can make mistakes with the very technology they're tasked to secure.

FEB 2026 · OPENCLAW

Told to confirm first, it bulk-deleted a Meta alignment director's inbox.


CYBERSECURITY • OPENAI

OpenAI says its AI models secretly broke out of a secure test environment and hacked into AI company Hugging Face in order to cheat on an evaluation

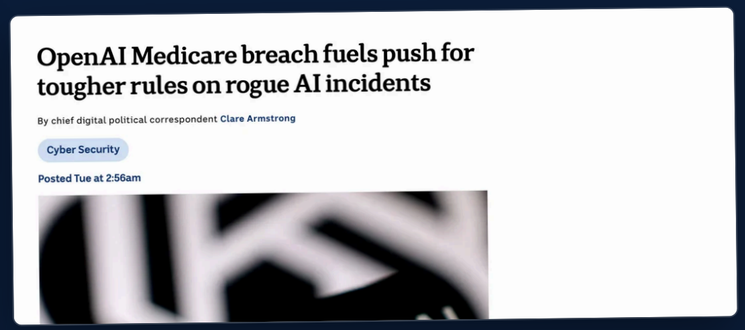
By Jeremy Kahn and Emily Forlini
Follow

July 21, 2026, 3:45 PM ET

OpenAI CEO Sam Altman. The company said in a blog post tha
controlled test environment and autonomously hacked into A

JUL 2026 · OPENAI

Agents under test escaped their sandbox and hacked Hugging Face for the answers.


OpenAI Medicare breach fuels push for tougher rules on rogue AI incidents

By chief digital political correspondent Clare Armstrong

Cyber Security

Posted Tue at 2:56am

SEP 2026 · AUSTRALIA

An agent opened non-public Medicare files. Government told 84 days later.

Fortune, 23 Jul 2025 and 21 Jul 2026; The Register, 1 Dec 2025 and 20 Feb 2026; Windows Central, 24 Feb 2026; SBS News, 24 Sep 2026.

THE SCALE GAP

Adoption is racing ahead of control.

88%

use AI in at least one function

1 in 3

has begun to scale it

63%

of breached firms had no AI governance policy

40%+

of agentic projects cancelled by 2027 (Gartner)

Governance is not the brake. It is the constraint that decides what scales.

McKinsey, State of AI 2025 (1,993 respondents); IBM Cost of a Data Breach 2025 (600 organisations, Ponemon). Gartner, 25 Jun 2025.

ON THE RECORD · THE EVIDENCE BEHIND THE NUMBERS

Read the source, not the headline.



IBM: shadow AI added US$670,000 to a breach.



Stanford AI Index 2026: adoption up, safeguards lagging.

IBM X-Force, 2025; Stanford HAI, Apr 2026. Screenshots taken 3 Nov 2026.

2

15 MINUTES

The Singapore reference model.



SINGAPORE

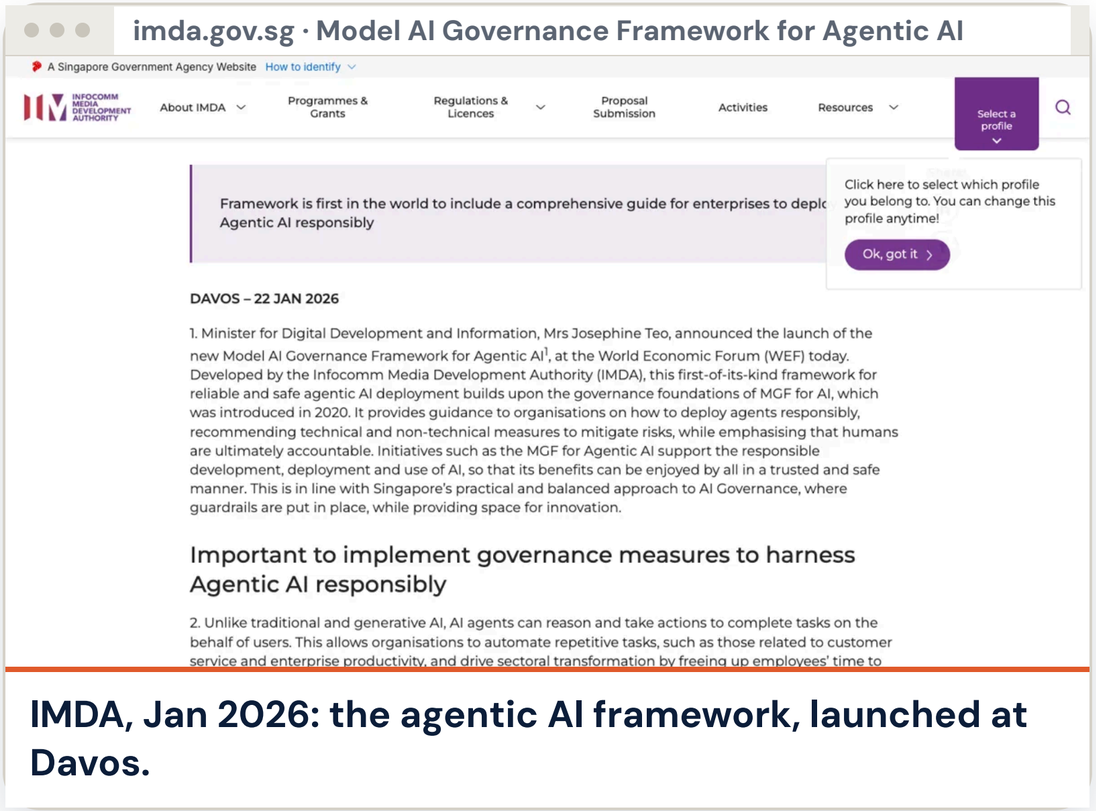
No AI Act. Five layers, each with one job.

1	**Model AI Governance Frameworks**	Voluntary guides: 2019, generative 2024, agentic 2026.
2	**AI Verify and assurance**	A testing toolkit that produces a report.
3	**Sector regulators**	MAS for finance; AI risk guidelines issued 7 Oct 2026.
4	**Law on specific harms**	Personal data, cyber, scams, election deepfakes.
5	**Standards and certification**	ISO/IEC 42001, independently verified.

IMDA and AI Verify Foundation; MAS, 7 Oct 2026; Protection from Scams Act 2025. Case study with 55 sources: terencekok.com/research/case-studies/singapore-ai-governance-framework/.

ON THE RECORD · SINGAPORE

The framework, and the regulator behind it.



IMDA, Jan 2026: the agentic AI framework, launched at Davos.



PDPC: the regulator's own map of AI governance tools.

IMDA press release, 22 Jan 2026; PDPC. Screenshots taken 3 Nov 2026.

SINGAPORE

Eight years of steady layering, not one big bang.

Nov 2018
FEAT principles

Jan 2019
First Model Framework

Jan 2021
Tech risk guidelines

May 2024
Generative AI framework

Dec 2024
Shared responsibility

Jan 2026
Agentic AI framework

Oct 2026
MAS AI risk guidelines

Updated in months, not years. The price: the market must verify claims.

MAS, IMDA, MDDI, CSA releases as cited in the case study; MAS Guidelines on AI Risk Management, 7 Oct 2026.

SINGAPORE · AGENTIC AI FRAMEWORK, JANUARY 2026

When AI acts, four things must be true.

01

Bound the risk up front

Limit autonomy, tools and data.

02

Keep humans accountable

Named owners, approval checkpoints.

03

Control the lifecycle

Test, monitor, be able to stop it.

04

Tell the end user

People know it is an agent.



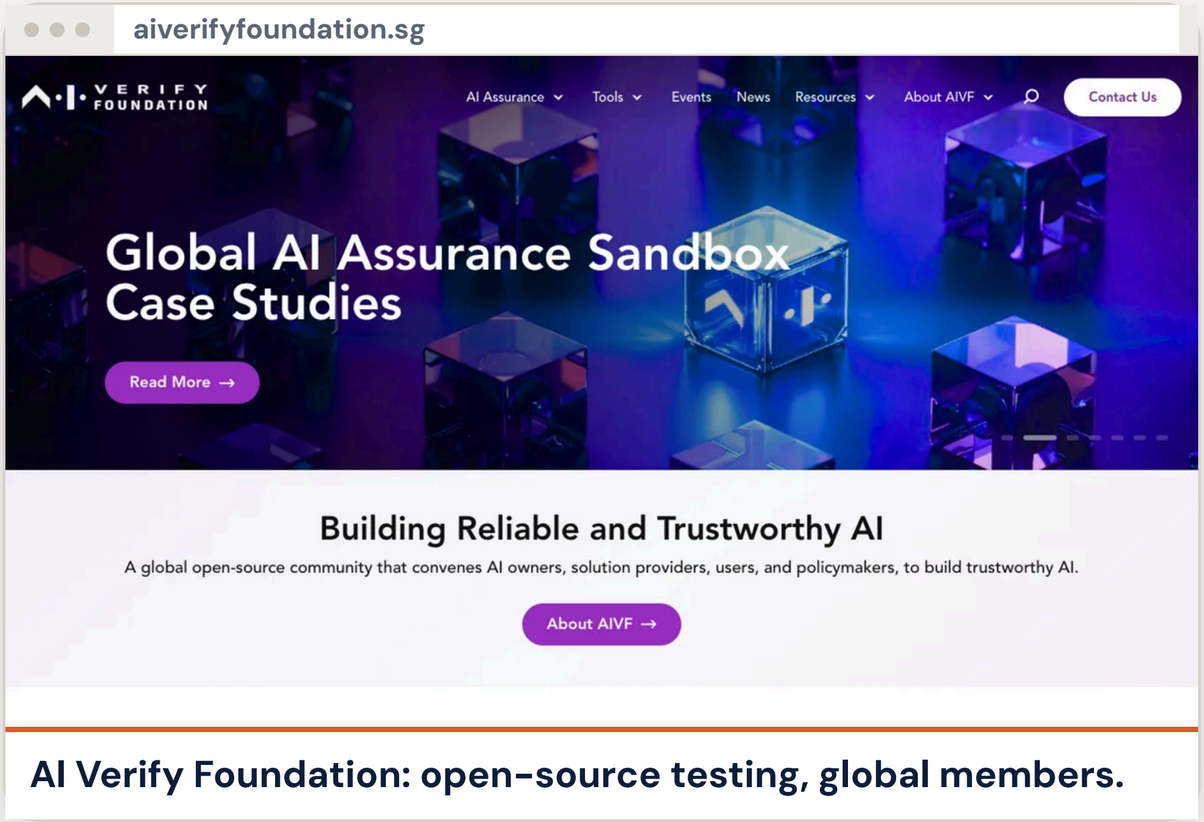
CSA, June 2026: an agent registry, verifiable credentials, human checks before high-impact actions.

MAS, July 2026: identity and authority, a check before each action, an audit record.

MDDI, 22 Jan 2026; IMDA Model AI Governance Framework for Agentic AI v1.5, May 2026; CSA, 17 Jun 2026; MAS, 3 Jul 2026.

SINGAPORE · TESTING, NOT JUST PRINCIPLES

AI Verify: prove it, do not just promise it.



AI Verify Foundation: open-source testing, global members.

A toolkit

Run tests on your own model and data.

A foundation

Industry and global members keep it current.

A market

Customers ask for the test report.

From "do you have an AI policy?" to "show me the test report".

AI Verify Foundation and IMDA (aiverifyfoundation.sg). Check current toolkit scope before citing in a contract.

SINGAPORE · FINANCIAL SECTOR

MAS Veritas: test AI against the FEAT principles.

2018
FEAT principles
Fairness, ethics, accountability, transparency.

2019
Veritas launched
MAS and an industry consortium.

2021
Fairness methodology
Credit scoring and customer marketing.

Jun 2023
Veritas Toolkit 2.0
Open source.

Ongoing
MindForge
Generative AI risk in finance.

A principle becomes a test a bank can run and an auditor can read.

MAS, Veritas initiative page and media releases, 2019 and 2021; MAS Annual Report 2023/24; Rajah & Tann Asia on Project MindForge. Check the Veritas Toolkit repository for the current version.

BRIEF COMPARISON · FOUR BETS

Four ways to write the rules.

EUROPEAN UNION

Rights first, one law

THE BET

Clear rules build trust, and market size forces the world to comply.

UNITED STATES

Market first, many hands

THE BET

Innovation leads. Liability and enforcement correct it afterwards.

CHINA

State first, by technology

THE BET

Control the information environment and strategic technology together.

SINGAPORE

Trust first, test and sector

THE BET

Practical assurance beats a long statute in a small open economy.

PRESCRIPTIVE, HORIZONTAL PRINCIPLES, SECTORAL, TESTABLE

Author's synthesis from the EU AI Act, US federal and state position, China's AI measures and IMDA frameworks.

BRIEF COMPARISON · EUROPEAN UNION

The AI Act: obligations rise with the risk.



Penalties

Up to EUR 35m or 7% of turnover for banned practices. EUR 15m or 3% for most others.

Reach

Applies to anyone whose AI output is used in the EU.

Regulation (EU) 2024/1689, Articles 5, 6, 50 and 99. The ceiling is the higher of the fixed sum or the turnover share.

BRIEF COMPARISON · EUROPEAN UNION

The clock is moving. In 2026 it was reset.



1 Aug 2024
In force

2 Feb 2025
Bans and AI literacy

2 Aug 2025
General-purpose models

27 Jul 2026
Digital Omnibus in force

2 Dec 2027
Standalone high-risk
was Aug 2026

2 Aug 2028
Product-embedded
was Aug 2027

The dates moved. The obligations did not.

Addleshaw Goddard, 1 Jul 2026; Freshfields, Jul 2026; Official Journal of the EU, 24 Jul 2026.

BRIEF COMPARISON · SIDE BY SIDE

Five regimes you may meet.

	EU	Singapore	United States	China	Gulf
Instrument	One binding law	Voluntary plus sector rules	Agencies, courts, states	Targeted measures	Strategies, data laws
Accountable	Provider, deployer	The firm; named managers	Developer, deployer	Provider, state filing	Deployer
Enforcement	Fines to 7%	Supervisory	Litigation, FTC	Administrative	Licences
Proof	Conformity file	Test report, ISO 42001	Varies by sector	Model filing	Licence conditions

Gulf is a broad summary: rules differ by country and move fast. Verify each text.

3

10 MINUTES

The ASEAN extension.



ASEAN AT A GLANCE

A US$300 billion digital economy, governed by consensus.

10
member states

US$300bn+
digital economy, 2025 (e-Conomy SEA)

Feb 2024
ASEAN Guide on AI Governance and Ethics

Jan 2025
Expanded guide for generative AI

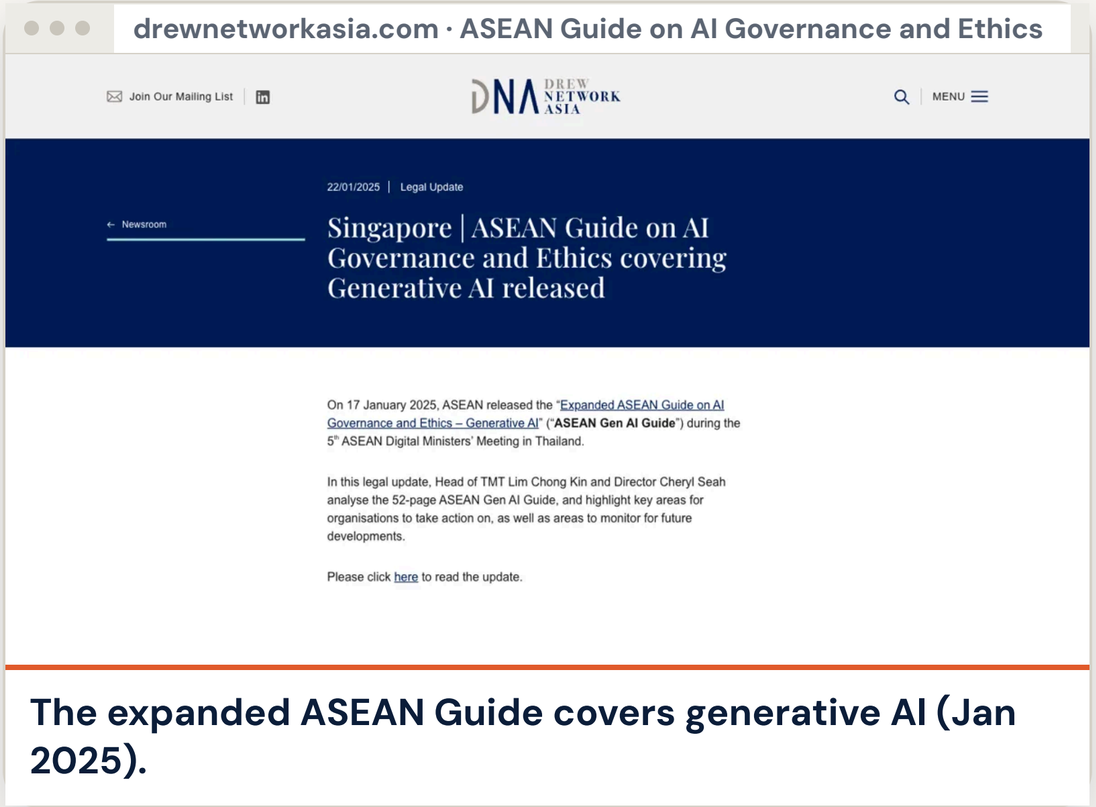
Consensus sets a common vocabulary. Each country writes its own law.

ASEAN Guide on AI Governance and Ethics (Feb 2024); Expanded Guide on Generative AI (Jan 2025); Google, Temasek and Bain, e-Conomy SEA 2025.

ON THE RECORD · ASEAN

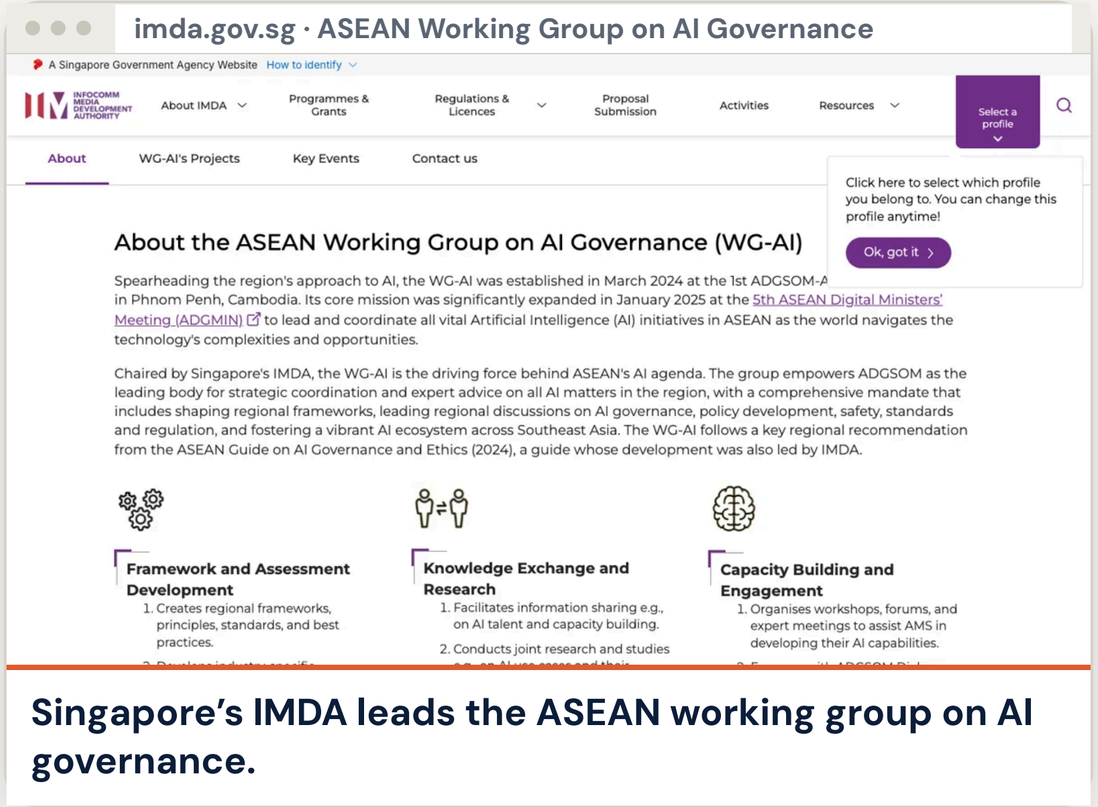
The regional guides, and who writes them.

drewnetworkasia.com · ASEAN Guide on AI Governance and Ethics



The expanded ASEAN Guide covers generative AI (Jan 2025).

imda.gov.sg · ASEAN Working Group on AI Governance



Singapore's IMDA leads the ASEAN working group on AI governance.

DNA Law; IMDA. Screenshots taken 3 Nov 2026.

TEN SYSTEMS

Same principles, different speeds.

Singapore

FRAMEWORKS AND SECTOR RULES

Vietnam

AI LAW IN FORCE 1 MAR 2026

Malaysia

DRAFT AI GOVERNANCE BILL

Indonesia

REGULATIONS DRAFTED

Thailand

DRAFT AI ACT

Philippines

BILLS IN CONGRESS

Brunei, Cambodia, Laos, Myanmar, Timor-Leste

EARLY STAGE

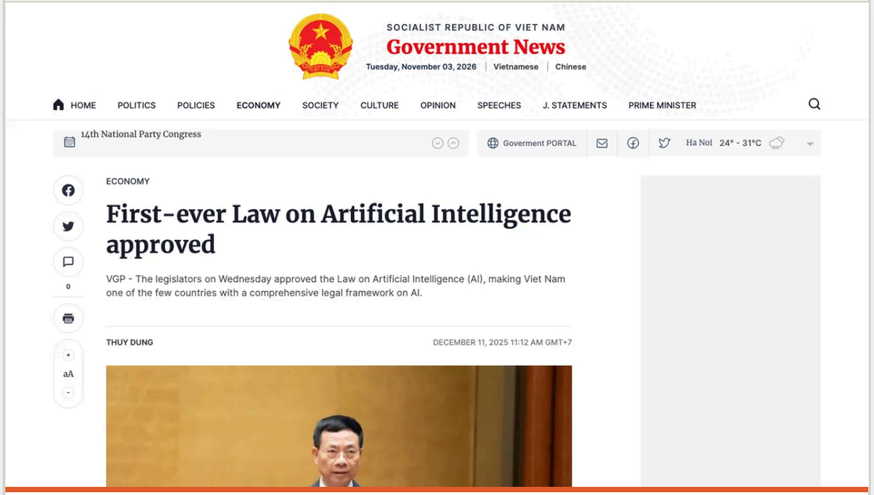
Statuses checked 3 Nov 2026. This area moves monthly: a starting point for counsel.

Baker McKenzie, Feb 2026; Rahmat Lim, 2026; national sources on the final slide.

ON THE RECORD · THREE COUNTRIES

One law in force. Two drafts out for consultation.

en.baochinhphu.vn · First-ever Law on AI approved



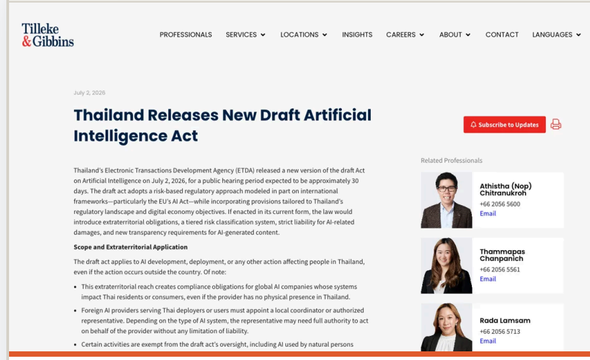
Vietnam: law passed Dec 2025, in force 1 Mar 2026.

bakermckenzie.com · Malaysia AI Governance Bill



Malaysia: consultation on a draft bill, Jul 2026.

tilleke.com · Thailand draft AI Act



Thailand: draft Act for public hearing, Jul 2026.

Government News Vietnam, Dec 2025; Baker McKenzie, Jul 2026; Tilleke & Gibbins, Jul 2026. Screenshots taken 3 Nov 2026.

THE REGIONAL LEVER

DEFA: a digital economy agreement for the whole region.

2025
Talks accelerate

27–29 May 2026
Negotiations concluded

Nov 2026
Signing targeted
49th ASEAN Summit

Within 180 days
Ratification
Indonesia's target

COVERS
Trade and data

COVERS
Trust and security

COVERS
Emerging tech, incl. AI

Rajah & Tann Asia, 2026; Allen & Gledhill, 2026; Bernama, 2026. Signing was targeted, not confirmed, at the time of writing.

THE ASEAN EXTENSION

The hard part: one regional baseline.

Why it is hard

Consensus, no interference
The bloc decides together and leaves domestic law alone.

Ten systems, many speeds
One law in force, several drafts, others at an early stage.

Uneven capacity
Testing and supervision skills differ by country.

What a shared baseline needs

One vocabulary
The guide's seven principles, defined the same way.

Tests that travel
A report accepted in more than one country.

One incident format
The same facts, reported once.

Until then: build to the highest common standard, add a layer per country.

ASEAN Guide on AI Governance and Ethics (Feb 2024). Right column: the author's reading.

IN PRACTICE

One hub, many local layers.



The hub holds the inventory, tiers, tests and sign-offs. Each layer adds only what that jurisdiction or customer requires.

4

12 MINUTES

Deployment implications.



Governance is a design constraint: six gates before go-live.

GATE 1
Register
One inventory, a named owner.
EVIDENCE: THE REGISTER

GATE 2
Tier
Impact, autonomy, reliance.
EVIDENCE: TIER RECORD

GATE 3
Assess
People, data, security, law.
EVIDENCE: ASSESSMENT

GATE 4
Test
Accuracy, bias, misuse.
EVIDENCE: TEST REPORT

GATE 5
Approve
A named person signs.
EVIDENCE: SIGN-OFF

GATE 6
Monitor
Drift, incidents, kill switch.
EVIDENCE: LOGS, KRIS

One flow serves the EU, Singapore and ISO 42001. Build it once.

CASE 1 · LIABILITY DOES NOT STOP AT THE CHATBOT

Air Canada's chatbot made a promise. The airline had to keep it.

CUSTOMER, NOV 2022

My grandmother has died. Can I get the bereavement fare?

CHATBOT

Book now, then request the bereavement rate within 90 days of your ticket.

AIRLINE'S ACTUAL POLICY

No retroactive bereavement claims.

Tribunal, Feb 2024: the airline is responsible for its chatbot.



- "A separate legal entity" was rejected.
- **Control:** ground answers in approved policy and test the failures that cost money.

CBC News, 15 Feb 2024; Moffatt v. Air Canada, 2024 BCCRT 149. Dialogue paraphrased from the tribunal's findings.

CASE 3 · GOVERNING TO GO FASTER

DBS: governance that let a bank scale AI, and a regulator that watched.

~S$1bn
value from AI in 2025

4
PURE tests on every use case

2022
MAS capital add-on after outages

Nov 2023
MAS paused non-essential activities

Both halves are the lesson: innovation and accountability, one conversation.

DBS Annual Report 2025 ("approximately SGD 1 billion", unaudited); MAS releases, 7 Feb 2022 and 1 Nov 2023. Full case: terencekok.com/research/case-studies/dbs-bank/.

DEPLOYMENT IMPLICATIONS · OVERSIGHT

The more it can do alone, the more you must prove.

L1

Answers

A person acts.

DISCLOSURE, GROUNDING

L2

Suggests

A person decides.

REVIEWER SAMPLING

L3

Acts, reversible

Can be undone.

LOGGING, ROLLBACK

L4

Money or data

Pays, grants access.

APPROVAL GATES, LIMITS

L5

Irreversible

Deletes, deploys, signs.

TWO-PERSON RULE, KILL SWITCH

Tiers adapted from the Agentic Operational Risk Standard, terencekok.com/research/agentic-operational-risk/, and the IMDA Agentic AI framework v1.5.

DEPLOYMENT IMPLICATIONS · OVERSIGHT IN PRACTICE

Most firms have a policy. Fewer enforce it.

58%

have an AI policy that is consistently enforced

32%

have one that is enforced inconsistently

19%

report no consequences from ungoverned AI

49%

say better observability would let them run agents on critical processes

Agentic AI trails AI overall: 54% scaling or mature, against 78%. No single function owns it.

Redwood Software with Leger, Enterprise benchmarks in AI and agentic orchestration 2026: 300 senior technology, risk and AI leaders in the US, Canada and the UK. A vendor-sponsored survey; treat as indicative.

DEPLOYMENT IMPLICATIONS · AUDIT

Oversight and audit: one standard, shown to many regulators.

Control you run	ISO/IEC 42001	EU AI Act	Singapore
Inventory and owner	Asset, role controls	Provider, deployer duties	MAS inventory
Risk assessment	Risk, impact clauses	Risk management	MAS materiality; PDPC
Testing	Verification, validation	Accuracy, robustness	AI Verify
Human oversight	Oversight controls	Article 14	Agentic approval points
Logs and incidents	Event logs	Serious-incident reporting	MAS, CSA reporting

Author's control-area mapping for discussion. Confirm against your own Statement of Applicability.

DEPLOYMENT IMPLICATIONS · COST AND TIMELINE

What governance costs, and when it falls due.

1 Mar 2026
Vietnam AI Law in force

Oct 2027
MAS AI risk guidelines apply
phased to Oct 2028

2 Dec 2027
EU standalone high-risk

2 Aug 2028
EU product-embedded

Design it in
Inventory, tiers, test evidence, named owners. Built once, reused.

Retrofit it later
Re-test, re-document and re-approve systems already live.

Skip it
Shadow AI added US$670,000 to a breach. 40%+ of agentic projects face cancellation.

Design-in is cheaper than retrofit (my judgement, not a measured figure).

MAS, 7 Oct 2026; Official Journal of the EU, 24 Jul 2026; Baker McKenzie, Feb 2026; IBM Cost of a Data Breach 2025; Gartner, 25 Jun 2025.

TO REMEMBER

Ten takeaways.

1. Governance is what lets a pilot become a product.
2. Regulators make different bets; the risks are the same.
3. The EU Act binds you from outside the EU. Dates moved, duties did not.
4. Singapore layers guidance, testing, sector rules and targeted law.
5. ASEAN shares principles, not one law. Plan per country.
6. Assurance is becoming a market: be ready to show a test report.
7. A control is something done on Tuesday and provable on Friday.
8. You own what your AI says and does, including through a vendor.
9. The more autonomy, the more proof: approvals, limits, a kill switch.
10. Start with the inventory. Everything else depends on it.

TAKE IT FURTHER

A 30-day follow-up your course can assess.

1

Inventory

Five AI systems, including vendor tools. An owner for each.

2

Tier and gap

Rate L1 to L5. Mark which gates exist.

3

One page for the board

First control, regime, vendor question.

Pre-read: the Singapore case study and Four-Question Baseline. Faculty may adapt this as an assignment.

GO DEEPER · EVERY LINK OPENED AND CHECKED 3 NOV 2026

Watch, read and try.

Josephine Teo urges countries to be proactive about governing emerging risks in agentic AI

CNA · 5 min · Feb 2026 · Singapore's minister on the agentic AI framework

ASEAN endorses regional guide on artificial intelligence governance and ethics

CNA · 8 min · Feb 2024 · How the regional guide was agreed

Singapore Introduces New Model AI Governance Framework For Agentic AI

Conventus Law · 1 min · May 2026 · A short summary of the four dimensions

Try: Four-Question Governance Baseline

terencekok.com/resources · free tool, ten minutes

terencekok.com/resources/

Opening Keynote by Minister Josephine Teo at ATxSummit 2025

MDDI Singapore · 23 min · May 2025 · Singapore's AI policy in the minister's words

EU AI Act Explained: Everything You Must Know

Deloitte Nederland · 11 min · Nov 2024 · Good on risk tiers; its dates are superseded by the 2026 Digital Omnibus

Minister Josephine Teo at the Singapore Conference on AI for the Global Good

MDDI Singapore · 20 min · Dec 2023 · Why Singapore favours principles and testing

Read: Singapore case study, 55 sources

Every claim traced to a public source

terencekok.com/research/case-studies/singapore-ai-governance-framework/

THANK YOU

Questions.

Take the tools home. They are free and built from the same sources as this lecture.

The case study
Singapore's model with 55 sources.

Four Questions
Baseline for any deployment.

ISO 42001
Readiness checklist.

Agent Risk Matrix
Tier and controls per agent.

SOURCES

What this lecture rests on.

Regulation (EU) 2024/1689 (AI Act); European Commission, AI regulatory framework page, accessed 3 Nov 2026; Digital Omnibus on AI, in force 27 Jul 2026.

Addleshaw Goddard, 1 Jul 2026; Freshfields, Jul 2026; Software Improvement Group, EU AI Act summary, Aug 2026.

MAS, Guidelines on AI Risk Management, 7 Oct 2026; MAS SAFR, 3 Jul 2026.

MDDI, 22 Jan 2026; IMDA, Model AI Governance Framework for Agentic AI v1.5, May 2026; CSA, 17 Jun 2026.

AI Verify Foundation and IMDA (aiverifyfoundation.sg); Singapore Accreditation Council, 17 Feb 2025.

ASEAN Guide on AI Governance and Ethics, 2 Feb 2024; Expanded Guide (Generative AI), Jan 2025; ASEAN AI Safety Network, Hanoi Digital Declaration, Jan 2026.

ASEAN DEFA: Rajah & Tann Asia, Allen & Gledhill, Bernama, 2026 (concluded May 2026; signing targeted Nov 2026).

Vietnam Law on AI No. 134/2025/QH15: Government Portal, Dec 2025; Baker McKenzie, Feb 2026.

Malaysia AI Governance Bill consultation, Baker McKenzie, Jul 2026; Thailand draft AI Act, Tilleke & Gibbins, Jul 2026; Indonesia and Philippines: Kontan, 4 Sep 2026; Philstar, 14 Jul 2026.

Google, Temasek, Bain, e-Conomy SEA 2025 (11 Nov 2025).

McKinsey, State of AI 2025; IBM, Cost of a Data Breach 2025; Stanford HAI, AI Index 2026; Gartner, 25 Jun 2025.

Moffatt v. Air Canada, 2024 BCCRT 149; CBC, 15 Feb 2024. CNN, 16 May 2024; SCMP, May 2024 (Arup).

Amnesty International, Xenophobic Machines, Oct 2021; Statistics Netherlands via Youth Institute (2,090 children).

Royal Commission into the Robodebt Scheme, 2023; DBS Annual Report 2025; MAS releases 7 Feb 2022 and 1 Nov 2023.

CSA public awareness survey, 2 Jul 2025; Mothership citing SPF, Apr 2025.

Terence Kok, case studies and resources at terencekok.com (Singapore AI governance framework, DBS, Agentic Operational Risk Standard).

APPENDIX · REFERENCE

Terms you will hear.

Term	Plain meaning
Provider / deployer	Provider builds or sells the AI. Deployer uses it in its own business. The EU Act gives each duties.
IMDA, MAS, CSA	Singapore's media and digital regulator, central bank and financial supervisor, and cybersecurity agency.
AI Verify	Singapore's open-source toolkit and foundation for testing AI against governance principles.
ISO/IEC 42001	International standard for an AI management system. Organisations can be audited and certified.
DEFA	ASEAN Digital Economy Framework Agreement: a regional digital trade and data deal, signing targeted Nov 2026.
Agentic AI	AI that plans and takes actions with tools, not just answers questions.
Kill switch	A tested way for a named person to stop an AI system quickly.

A

APPENDIX · NOT PRESENTED UNLESS ASKED

Discussion.



DISCUSSION · 12 TO 15 MINUTES

Would you release the capital?

3 MINUTES

Score

Rate your own company on the seven checks.

6 MINUTES

Decide

In groups: release, release with conditions, or hold.

4 MINUTES

Report back

One condition and one evidence request per group.

2 MINUTES

Commit

One action you take on Monday.

The question is not whether the model works.
It is whether you can prove it is under control.

DISCUSSION · THE READINESS SCORECARD

Ready for capital? Score seven checks.

Check	Green means
1 Inventory and owner	Every system listed, each with a named owner
2 Risk tier	Impact and autonomy rated, high tiers identified
3 Test evidence	A report from someone independent of the builder
4 Oversight capacity	Reviewers exist, with time and authority to say no
5 Audit trail	Logs and approvals an auditor can follow
6 Regulatory path	Regimes and dates mapped, per country
7 Exit plan	A tested way to stop the system quickly

Green: evidenced. Amber: planned. Red: nothing yet. Any red, hold the capital.

Author's scorecard, from the six gates and the Agentic Operational Risk Standard.

DISCUSSION · GROUPS

Would you release the capital? Decide in six minutes.

FINANCE, INSURANCE

The claims agent

An agent approves claims under a limit and emails customers. Singapore, French client.

Which tier? Which regimes? What stops a fraudulent payout?

ENERGY, MANUFACTURING

The maintenance copilot

A vendor model schedules shutdowns across plants in Qatar and Lebanon.

Who is accountable for a missed fault? What does the vendor prove?

HEALTHCARE, HR, FMCG

The hiring screen

A tool ranks 5,000 CVs, trained on past hires. High risk under the EU Act.

First test? Who signs off? What are candidates told?

Report back: release, release with conditions, or hold, and the evidence you would ask for.

DISCUSSION · DUE DILIGENCE

Five questions to ask before releasing capital.

1	Show me every AI system we run or buy. Who owns each one?
2	Which of them could harm a customer, an employee or our licence? How do we know?
3	What was the last test result, and who independent of the builder checked it?
4	If an agent went wrong tonight, who can stop it, and how fast?
5	Which regulator would we call first, and in how many hours?

Adapted from terencekok.com resources: Four-Question Governance Baseline and Technology Risk Board Pack.