JAMES A LANG

BOARD BRIEFING · NOVEMBER 2026

FOURTEEN PRINCIPLES FOR GOVERNING AI

The short version for the boards of critical national infrastructure



James A Lang
AI non-executive director
www.jamesalang.com

THE SHORT VERSION

FOURTEEN PRINCIPLES. ONE AIM.

Protect the investment the board makes in AI, and the value that investment returns.

A board has three jobs with AI: protect the value, set the appetite for risk, and keep the power to stop. It cannot do any of them on assurances. It needs evidence.

> **The division is simple.** The board owns the principles and sets the appetite. The AI Office makes sure every principle is carried out, and tells the board when one is not. The executive does the work.

FOUR GROUPS, SO THEY ARE EASY TO HOLD IN MIND

VALUE	We know what AI is worth, and we get it	**Principles 1 and 2**
RELIABLE	We know who decides, and within what limits	**Principles 3 to 6**
RESILIENT	We can withstand attack and failure, and we can stop	**Principles 7 to 10**
RESPONSIBLE	People stay in control and are treated properly	**Principles 11 to 14**

WHAT MAKES THESE DIFFERENT

- **Value comes first.** Principle 1 asks what the board expected, what it has received, and the rule for stopping what does not pay.
- **A person stays in control.** Principle 11 keeps a named person in control of every AI system. No one can override the head of safety's signature.
- **Written for infrastructure.** Where AI touches physical assets, a wrong answer is not a bad report. It is a pressure, a dose, a flow or a load.

WITH THE TEST A BOARD CAN APPLY TO EACH

THE FOURTEEN ON ONE PAGE

No.	Principle	The board's test
	VALUE We know what AI is worth, and we get it	
1	**Protect the investment and realise its value**	For each AI use, what did we expect, what have we received, and what is the rule for stopping it if it does not pay?
2	**Set strategic goals tied to the business**	Is the value of each use stated as money, time or learning, with the evidence behind it?
	RELIABLE We know who decides, and within what limits	
3	**Set the appetite, and assess impact and risk**	Has the board set its appetite in numbers, and does each risk carry a likelihood, a loss range and an owner?
4	**Fix board accountability and management responsibility**	Who, by name, can switch a system off today, and when was that last tried?
5	**Empower an independent, cross-functional review body**	Has it ever said no, and was the decision respected?
6	**Review systems, policy and governance regularly**	What changed since the last release decision, and who decided it was still acceptable?
	RESILIENT We can withstand attack and failure, and we can stop	
7	**Watch the regulatory and geopolitical environment**	Can the board name the three obligations that would stop a deployment?
8	**Audit and measure AI use continually**	Is there one register, and when was it last checked against what is running?
9	**Build security in from the start**	When was each service last tested from the attacker's side?
10	**Test, and remove from use if harm is found**	Which threshold, if breached, stops the system without further approval?
	RESPONSIBLE People stay in control and are treated properly	
11	**Keep AI human-centred and safe**	For each AI system, who is the person in control, and has the head of safety signed?
12	**Validate, document and secure data**	Has the head of data signed that the data is fit and lawful for each service?
13	**Comply with privacy requirements**	Is there an assessment for every AI use that touches personal data?
14	**Train people to use AI well**	Can a control room operator say when they would distrust the system?

Twelve of the fourteen take the Institute of Directors business paper AI Governance in the Boardroom as their reference point. Principles 1 and 11 are original to this set. The full white paper is published by Velinor.

WHO MAKES IT HAPPEN

PRINCIPLES NEED AN OWNER: THE AI OFFICE

Fourteen principles on paper change nothing. The AI Office is the small central team that makes sure each one is carried out.

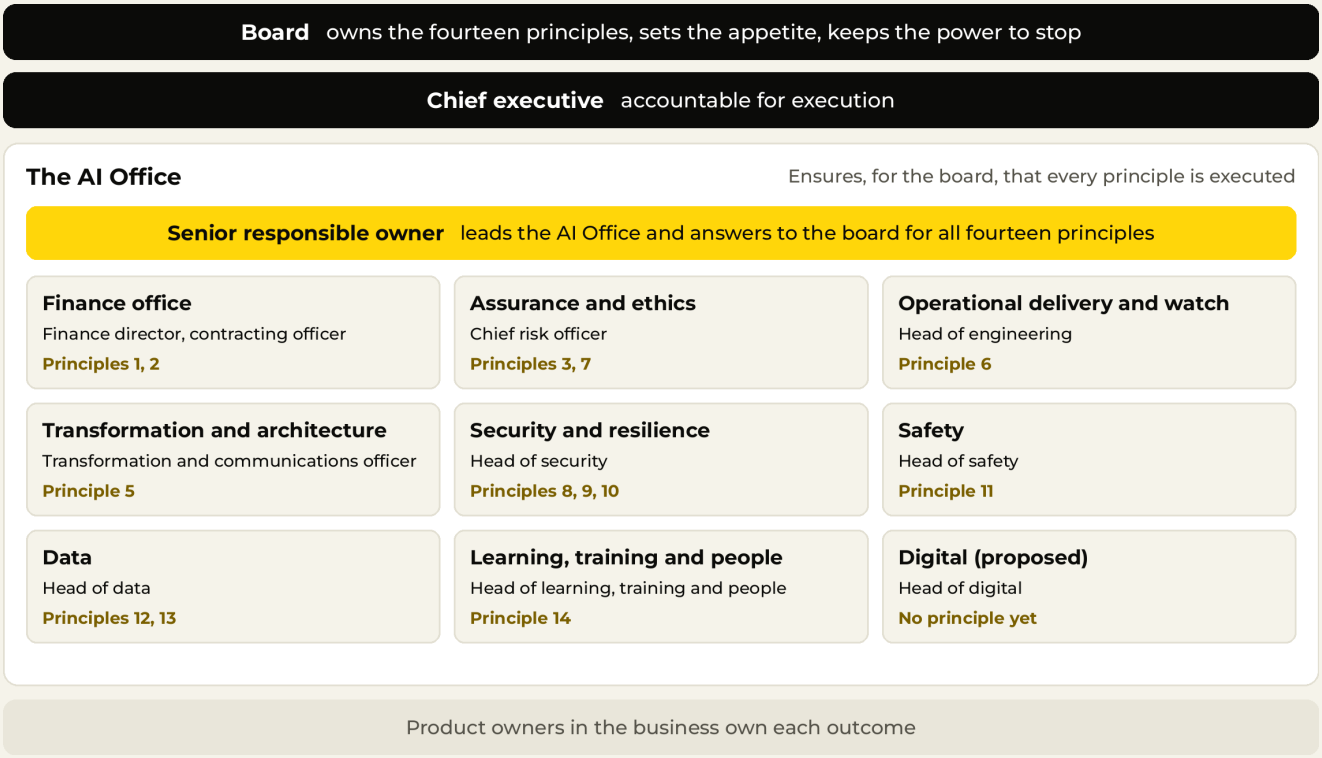

- **One senior responsible owner** leads the Office and answers to the board for all fourteen.
- **Every principle has a named lead.** The heads of data, security, safety and engineering sit inside the Office.
- **The board keeps two things.** The power to stop a system, and the duty to set the appetite.
- **Evidence, every quarter.** Each piece of assurance work leaves a product, and each product becomes a line in the board report.

THE FIRST 90 DAYS

START THIS QUARTER

A board can be working to all fourteen principles within one quarter. Know what exists, decide who decides, then test.

DAYS 1 TO 30

SEE IT

Adopt the principles. Commission one register of the AI in use. Name who can stop each system.

DAYS 31 TO 60

DECIDE IT

Set up the AI Office. Agree the appetite in numbers, and the rule for stopping a use that does not pay.

DAYS 61 TO 90

PROVE IT

Run one service through the full cycle, including a rehearsed stop. Receive the first board report built from the records.

THREE QUESTIONS FOR YOUR NEXT BOARD MEETING

1. Who, by name, can switch an AI system off today, and when was that last tried?
2. For each AI use, what did we expect, and what have we received?
3. Is there one register of the AI we run, and when was it last checked against what is running?

IF THE BOARD CANNOT ANSWER ALL THREE, SET UP YOUR AI OFFICE.

Find out how. Book a briefing at **www.jamesalang.com**, or scan the code on the back page.

NEXT STEP

SET UP YOUR AI OFFICE

The principles tell the board what to ask for.
The AI Office is how it gets it.
Find out how to set up yours.

Book a briefing

Scan the code, or visit
www.jamesalang.com

James A Lang
AI non-executive director
www.jamesalang.com

© James A Lang 2026. This paper is general guidance and is not legal advice.