Velinor

EXECUTIVE WHITE PAPER · NOVEMBER 2026

Protect the value of AI, from design to operation

How one small, central team gives boards AI they can trust, and the evidence to prove it.

VALUE RELIABLE RESILIENT RESPONSIBLE



James A Lang
Founder and CEO, Velinor
velinor.io



About this paper

A board-level guide to governing AI so that it protects value, not just reduces risk

This paper is for chairs, non-executive directors and executive teams who are accountable for AI they did not build and cannot see inside. It sets out the AI Office: what it is, how it works and what a board should ask for.

Start with the executive summary. Every section then opens with its finding as the headline and closes with what the board should do. The figures in section 07 come from an illustrative scenario and describe no real organisation.

Contents

	Executive summary	3
01	The problem	5
02	The AI Office	5
03	The standard	7
04	Organisation	7
05	Lifecycle	9
06	Decision rights	10
07	The investment	12
08	Trust	14
09	Blindspots	15
10	Incidents	16
11	Where it fits	18
12	Getting started	19
13	Board questions	20
14	Success	21
	How Velinor can help	22

Executive summary

AI is now a leadership responsibility. The AI Office makes it one that leaders can carry.

Ungoverned AI leaves leaders exposed. Boards are approving AI investment they cannot see inside, run by suppliers they do not control, and supporting decisions nobody has clearly approved. The AI Office is the small, central team that closes that gap, from the first design to the day a system is retired.

Five messages

1. **AI programmes rarely fail at the technology**

   They stall when the tool is switched on but the work does not change. Whether the work changes is the single variable that decides the return.

2. **A single number is not a business case**

   A board that approves AI spend on one benefit figure has accepted a risk it has not seen. Risk should be stated as an event, a likelihood and a range of loss in pounds.

3. **One small team holds four things**

   The value, the standard, the picture and the stop button. It works through existing executives and does not duplicate them.

4. **Accountability stays inside the organisation**

   Every AI-supported decision carries three names: who has authority, who is responsible and who is accountable. A supplier can run a system; it cannot hold the accountability.

5. **The result is clarity, control and confidence**

   One picture of the AI you run, named authority with the power to stop, and evidence rather than assertion that the investment is protected.

Three things to do this quarter

1. **Ask for the three answers.** The full AI register, a named owner for each AI-supported decision, and the likelihood of loss on the AI budget.
2. **Name the people.** A Senior Responsible Owner, a non-executive lead for AI, and who holds the stop button.
3. **Set the first gate.** Fix the date of the first value gate, and release no further AI capital without evidence.

Governance here is a shield, not a brake. Every control exists to protect the capital, the intent and the future the organisation has committed to.



01 · The problem

Most boards cannot yet answer three basic questions about their AI

AI now enters an organisation through many doors: funded projects, features switched on inside supplier platforms, and staff using public tools to get the job done. Each door adds value. Each also adds decisions that nobody has clearly approved.

Three questions most boards cannot answer today

1. What AI do we run, including the tools nobody sanctioned?
2. Who is accountable for each decision that AI supports?
3. How likely are we to lose money on our AI investment, and by how much?

The usual instruments do not help. Most AI business cases state one benefit figure and one cost figure. Most risk registers state red, amber or green. Neither tells a board how likely the investment is to lose money, or how big the loss could be.

Meanwhile accountability drifts. A supplier runs the model. A platform team runs the infrastructure. An operator acts on the output. When something goes wrong, each can point to the other. Regulators, shareholders and the public will not accept that answer.

The decision is the unit of accountability. Every AI-supported decision needs a named owner before it goes live.

For the board Ask management for three answers: the full AI register, a named owner for each AI-supported decision, and the likelihood of loss on the AI budget, stated as a number.

02 · The AI Office

One small team holds the value, the standard, the picture and the stop button

The AI Office is the central team that protects the value of the AI capability across its whole life: design, develop, deploy and operate. It owns the governance plan, runs it and answers for it. One message goes to everyone in the organisation: if it involves AI, the AI Office knows about it.

Illustration · The AI capability the AI Office protects: people, data, systems and suppliers

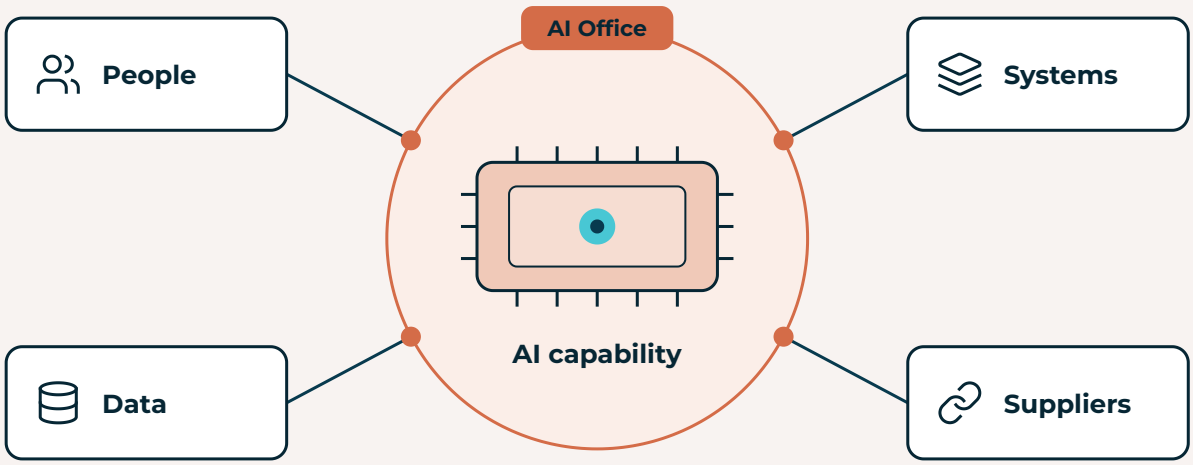

The mandate: four things the AI Office holds

The value

Track the return on the AI investment and report the likelihood of loss, in numbers.

The standard

No AI solution enters service unless it is reliable, resilient and responsible.

The picture

Keep one register of AI in use, suppliers and unsanctioned tools included.

The stop button

Hold delegated authority to halt a release or withdraw a solution from service.

What it is not

- **Not a new bureaucracy.** It is small and works through the existing risk, security, finance and operations functions. It does not duplicate them.
- **Not the owner of outcomes.** Product owners in the business own each outcome. The AI Office makes sure they can deliver it safely.
- **Not a brake.** Governance here is a shield. Every control exists to protect the capital, the intent and the future the organisation has committed to.
- **Not the EU AI Office.** Here the term means the organisation's own team, not the European Commission body that oversees the EU AI Act.

For the board Confirm who holds the stop button today. If the answer is a committee, or nobody, that is the first gap to close.

03 · The standard

Every AI solution meets one standard: reliable, resilient and responsible

The AI Office holds every AI solution to one standard, Velinor's R³AI: reliable, resilient and responsible. It is a minimum condition of deployment, kept for the whole life of the system. It is not an aspiration.

R³AI: what each standard means and how the board knows

Standard	What it means	How the board knows
Reliable	The AI does what it was built to do, to a stated level of performance, on this organisation's data.	Performance against agreed thresholds, tested before release and monitored in service.
Resilient	The service keeps running when the AI fails, is attacked or is withdrawn. It can be paused, isolated or rolled back safely.	Tested fallback, recovery times met, and the integrity of models and data evidenced.
Responsible	Its use is lawful, fair, explainable and owned by a named person.	A named owner, a recorded decision and an impact assessment for every material use.

The responsible test matters most because it catches the failure nobody sees. An AI that is biased or harmful while working exactly as designed has broken nothing technical. Only a standard that asks "is this right to use?" will find it.

For the board Adopt one standard for every AI release, and ask to see the evidence against it, not a statement that it was met.

04 · Organisation

The AI Office works through existing executives, not around them

The AI Office sits under the chief executive and is led by a Senior Responsible Owner, usually the executive who owns service performance. It has nine functions. The heads of data, security, safety and engineering sit inside it, and no other function can override the head of safety's signature.

Figure 1. The AI Office, who sits in it and who it answers to

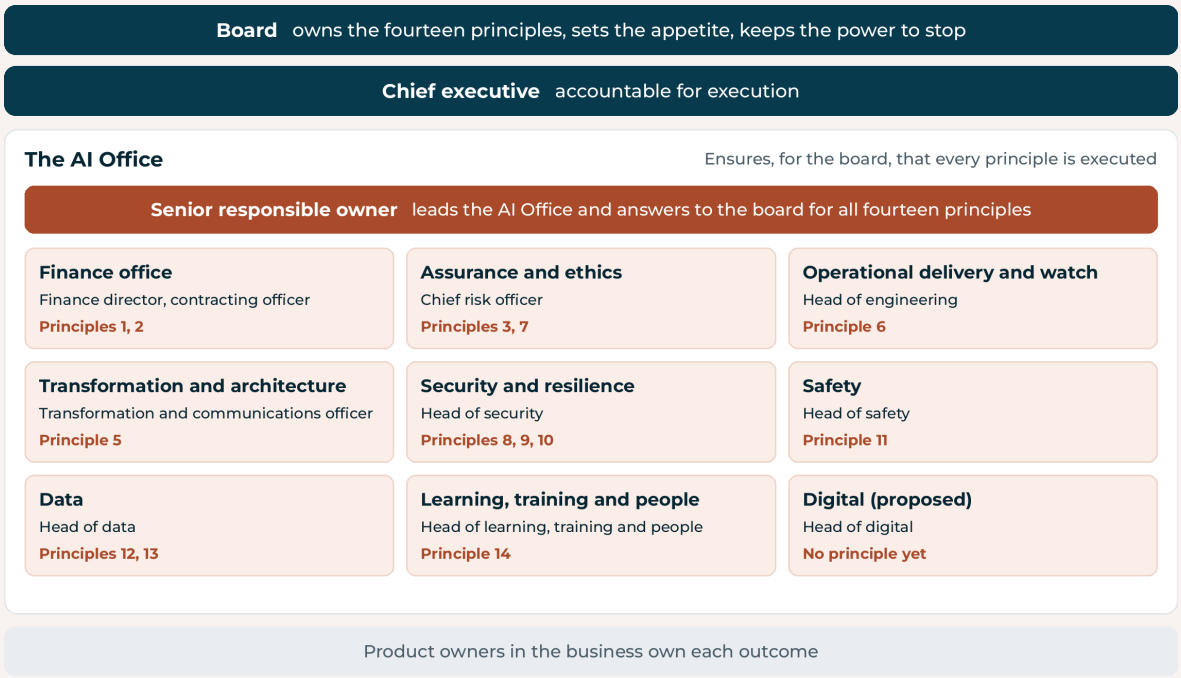

The nine functions

Function	What it does	Led by	Principles
Finance office	Runs the value-at-risk model, the value gates, benefit tracking, cost control and all AI contracting.	Finance director, with the contracting officer	1, 2
Assurance and ethics	Holds the standard, release approvals and impact assessments, and watches regulation.	Chief risk officer	3, 7
Security and resilience	Keeps the register complete. Security assurance, adversarial and containment testing, threat assessment and incident response for AI.	Head of security	8, 9, 10
Safety	Accepts AI into the safety case and keeps a named person in control of every system.	Head of safety	11
Data	Signs that the data is owned, lawful, of known origin and fit for use. Owns privacy.	Head of data	12, 13
Operational delivery and watch	Runs AI securely in service: monitoring, drift, retraining and retirement. Answers for outcomes in operation.	Head of engineering, with the operations director	6
Transformation and architecture	Holds one enterprise architecture for AI and leads the change that redesigns the work around it.	Transformation and communications officer	5
Learning, training and people	Workforce plan, role-based training, competence and human oversight.	Head of learning, training and people	14
Digital (proposed)	Owns the platforms AI is built and run on.	Head of digital	None yet

The Senior Responsible Owner leads the Office and answers for principle 4. The head of safety matters most where AI touches physical things: a treatment works, a network, a plant. There a wrong answer is not a bad report. It is a pressure, a dose, a flow or a load. No other function in the AI Office can override the head of safety's signature.

For the board Name a Senior Responsible Owner and a non-executive lead for AI, and confirm in writing who can stop a release.

05 · Lifecycle

Assurance is built in from design, not bolted on at release

Security, data, engineering and safety are not consulted at the end. They are built in from the start, not bolted on at the end. The AI Office works to the four stages in the NCSC's guidelines for secure AI system development, and checks all three standards at each one.

Illustration · Design, develop, deploy, operate: one loop, assured at every stage

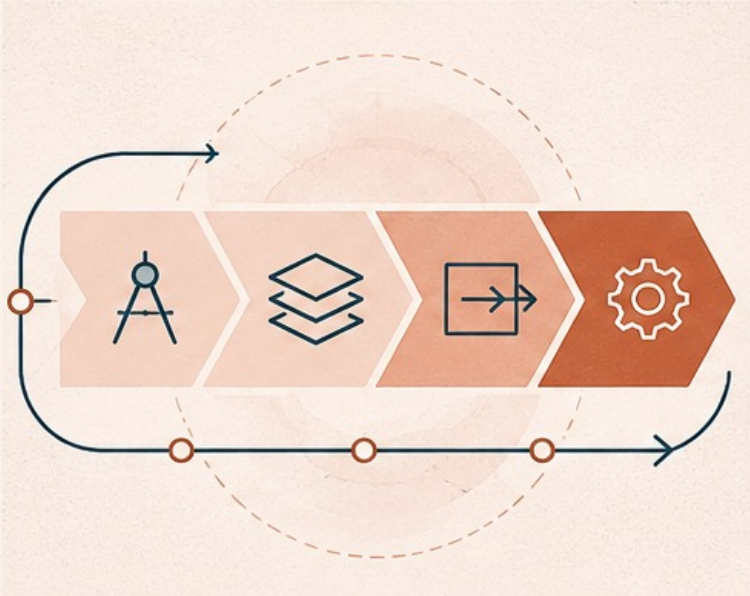

What the AI Office assures at each stage

Stage	Reliable	Resilient	Responsible
Design	Named outcome, measured baseline, performance threshold	Threat model; fallback designed in	Purpose approved; impact assessment; oversight mode set
Develop	Tested on the organisation's own data; data quality evidenced	Supply chain checked; models, data and code signed and versioned	Fairness tested; limits documented
Deploy	Release thresholds met; users trained for the task	Adversarial and containment tests passed; incident plan ready	Named owner signs; people affected are told

Stage	Reliable	Resilient	Responsible
Operate	Performance, drift and benefits monitored	Logging, patching, threat intelligence, exercised response	Complaints and overrides tracked; retired cleanly

For the board Require data, security, engineering and safety sign-off at the design stage, not at release.

06 · Decision rights

Every AI-supported decision carries three names, and no one releases AI alone

Decisions are no longer moments that executives make. They are architectures that executives are accountable for. So every AI-supported decision is recorded with three names against it. No AI solution enters service without all three. Where a supplier runs the system, accountability stays inside the organisation.

Authority, responsibility and accountability

Term	Meaning	Usually held by
Authority	The right to approve, change or stop the AI solution	Senior Responsible Owner
Responsibility	The duty to build, run and monitor it day to day	Product owner and Head of AI
Accountability	Answerable to the board and the regulator for the outcome	A named executive director

The sign-off chain

No single function can release an AI solution alone. Each signature is evidence that the board and the regulator can later see. In a critical national infrastructure operator the chain runs in this order:

1. Product owner: the outcome, baseline and users are ready.
2. AI Office, assurance and ethics: the standard is met and the impact assessment is complete.
3. Head of data: the data is owned, lawful, of known origin and fit for this use.
4. Head of security: the security assurance statement is accepted.
5. Head of engineering, where AI touches the engineered asset: the change has passed management of change.
6. Head of safety, where AI could affect safety: the AI is accepted into the safety case.
7. Chief risk officer: the residual risk is inside the board's appetite.

8. Senior Responsible Owner: release is authorised.

A refusal at data, security, engineering, safety or risk stops the release. Only the board can overrule one, and it does so on the record. Smaller organisations merge roles; the principle that no one signs alone stays.

Human oversight, set by consequence

Every solution is placed in one of three oversight modes before release. The mode is set by the consequence of error, not by what the technology can do.

Illustration · A trained person reviews the decision; the mode is set by the consequence of error

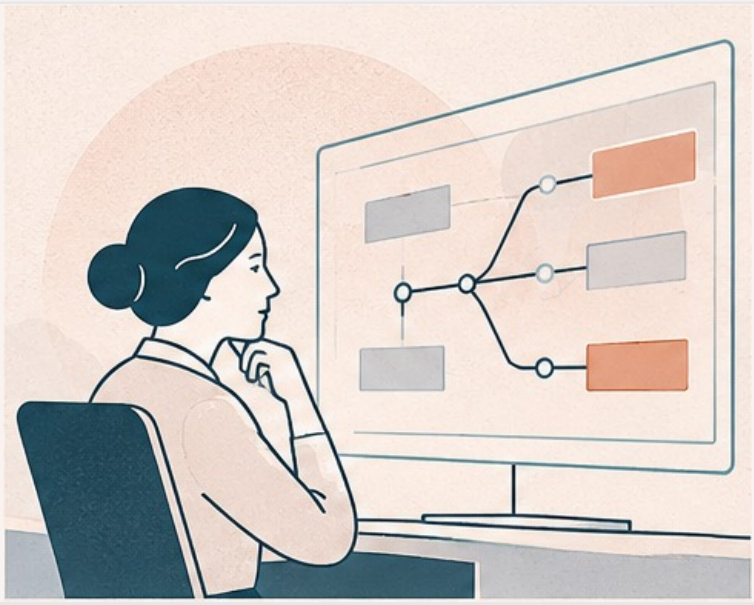

Oversight modes

Mode	The person's role	Used where
Human decides	AI advises. A trained person makes and owns the decision.	Safety-critical and service-critical decisions
Human supervises	AI acts inside limits. A person monitors and can stop it.	Reversible operational decisions with tested limits
Human audits	AI acts. People review samples and exceptions afterwards.	Low-consequence, high-volume tasks

No solution affecting an essential service runs in the third mode. And AI never replaces an engineered safety function: it may advise on protective limits, but it does not hold them.

For the board Approve no AI release without three names against it and a recorded oversight mode.

07 · The investment

Whether the work changes decides the return, and that can be measured

The finance office applies the measurement approach set out by Douglas Hubbard in *The Failure of Risk Management*. Applying it to AI is Velinor's extension, not the author's claim. Four rules follow.

1. **One language of risk.** Every risk is an event, a likelihood over a stated period and a range of loss in pounds. No colours without numbers behind them.
2. **Ranges, not points.** Every estimate is a 90% range from people trained to calibrate their estimates.
3. **Model the decision.** A Monte Carlo simulation combines the ranges and shows the spread of outcomes.
4. **Measure what matters.** Measurement effort goes only to the inputs that could change the decision.

Value gates

Capital is released in tranches. Each gate is passed on evidence, and the model is rerun with that evidence before the next tranche. The percentages are cumulative, and a starting point for agreement with the finance director, not a rule.

Figure 2. Value gates: cumulative capital released, and the evidence to pass

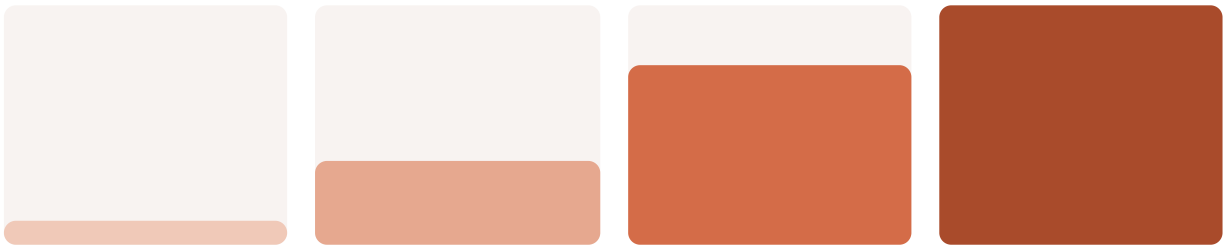

G0 Commit
Up to 10% released

Named outcome and owner, calibrated ranges, measured baseline, model run.

G1 Prove
Up to 35% released

One use case has changed how the work is done, and its measure has moved.

G2 Scale
Up to 75% released

Benefit range narrowed by live data; capability scored and no element weak.

G3 Sustain
Up to 100% released

Benefits tracked against forecast for two quarters; run cost inside its range.

Illustrative only

A worked example: £45 million of AI infrastructure capital over three years at a critical national infrastructure operator. Every input is invented to show the method; none comes from a client. 200,000 simulations, seven years, 3.5% discount rate. Likelihoods are on a scale of 0 to 1.

Measure	Without an AI Office	With an AI Office
Likelihood the investment loses money	0.67	0.46
Median net present value	−£42.6m	+£6.5m
90% range of net present value	−£95m to +£121m	−£78m to +£148m
Likelihood of at least £10m of AI incident losses in seven years	0.26	0.09

The governance effect is an assumption in this model, not a measurement: it assumes the AI Office raises the likelihood that the work actually changes from 0.45 to 0.70, and halves the yearly likelihood of a costly incident. Year-one forecast scoring is how that assumption gets tested.

Three findings matter to a board. All three come from the illustrative example above.

0.57

Changing the work decides it

The investment breaks even only if the likelihood that AI changes how the work is done is above about 0.57. Ask management to evidence it.

£5.5m

Governance pays for itself

It costs about £5.5m over seven years in present value and adds an expected £31.7m, on these ranges. The plan must prove that.

0.46

Still above appetite

Even with the AI Office, a 0.46 likelihood of loss is above an illustrative appetite of 0.35 at G2. Capital to scale waits for live evidence.

The model also prices better information. In the example, knowing early whether the work will change is worth up to £18.6m, and knowing the true size of the yearly benefit up to £10.5m. Better estimates of capital overrun and run cost are worth almost nothing to the decision. So the first measurement spend goes on an honest early test on one named outcome, and a live pilot to narrow the benefit range.

For the board Release capital in tranches, on evidence. Ask for the likelihood that the work changes, and the evidence behind it.

08 · Trust

Trust is given by others. Trustworthiness has to be evidenced

Trustworthy is not the same as trusted. Trust is what others give. Trustworthiness is what the organisation can evidence. Following the philosopher Onora O'Neill, it has three parts: honesty, reliability and competence.

The AI Office measures it with the Velinor Trust Framework (VTF). The framework has four pillars, 12 measurement areas and 40 interview questions, each scored on a five-point maturity scale.

Illustration · Four pillars, built in order, each resting on the one before



The Velinor Trust Framework: four pillars

Pillar	The question it answers
Trusted Purpose	Why does AI exist here, and what must it never compromise?
Trusted Leadership	Who decides, who is accountable, and how do we correct course?
Trusted Capabilities	Are our people, data, systems and suppliers fit to deliver it safely?
Trusted Outcomes	Are the results reliable, fair, explainable and of real benefit?

The order matters. Capability built before the purpose is clear delivers the wrong thing well. Outcomes claimed without accountable leadership cannot be defended. In the first 30 days the AI Office scores the organisation from 1 (Not Started) to 5 (Leading), and reports the lowest pillar and lowest area beside the overall score, because an average can hide a weak area. That baseline is what the board measures progress against.

For the board Commission a trust baseline in the first 30 days, and track the lowest score, not the average.

09 · Blindspots

The risks that hurt most are the ones no standard describes yet

Standards describe risks someone has already met. The events that damage an organisation most are the ones no standard yet describes. The AI Office manages three kinds of risk.

Illustration · Layers of defence around the essential service, with one signal watched

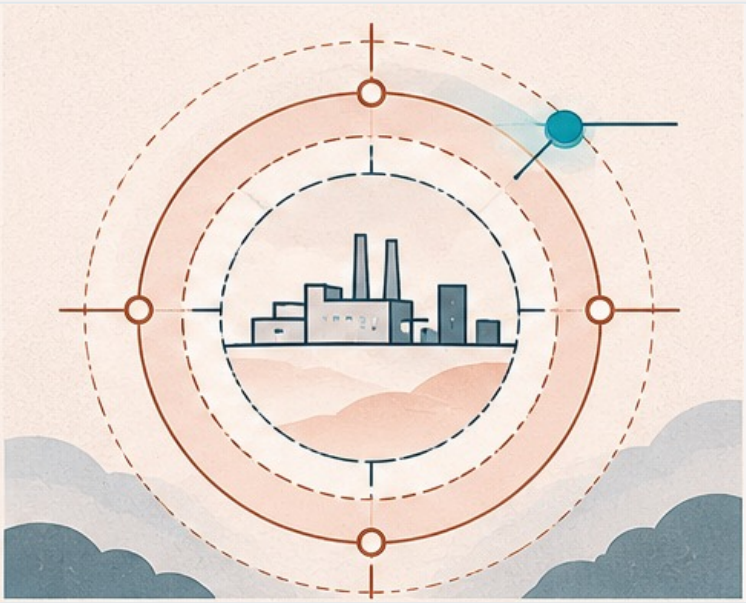

Three kinds of risk

Kind	What it is	How the AI Office manages it
Known	Risks written into standards, regulation and the risk register	Controls, testing and assurance
Emerging	Threats and failures seen elsewhere, not yet here	A standing blindspot watch: open-source intelligence, graded and acted on

Kind	What it is	How the AI Office manages it
Unforeseen	An event outside what anyone planned for	The ability to adapt fast: exposure known quickly, rehearsed response, tested fallback

Velinor's AIBlindspot framework maps where AI programmes fail in ways their leaders did not see coming, across the four lifecycle stages. The AI Office reviews every blindspot category at least once a quarter. A category with no findings is reported as "looked and found nothing", never left blank. A "don't know" is not a low risk. It is a task to find out.

When evidence changes, the likelihoods in the risk model are recalculated and the board sees the movement. Adaptability is measured, not assumed. Each quarter the board sees three numbers:

- **Time to know:** how long it takes to establish exposure to a new threat.
- **Time to act:** how long it takes to contain, patch or withdraw.
- **Time to learn:** how long before the lesson is in the controls and the model.

Resilience buys the right to be wrong about what was coming, and still be standing.

For the board Ask each quarter what happened to a peer that could happen here, and how long it took us to know.

10 · Incidents

Every AI solution will fail. What matters is how fast you know, contain and learn

Every AI solution will fail at some point. What matters is whether the organisation knows quickly, contains the harm and learns. The AI Office runs AI incidents on V-AIM, Velinor's six-stage incident model.

Figure 3. V-AIM: six stages from first signal to learning

1 Prepare

Register, rollback, named command roles and a tabletop exercise in place before any incident.

2 Detect

Decide whether something is actually wrong, and declare an incident if it credibly might be.

3 Contain

Preserve the evidence first, then pause, isolate or move to human control.

4 Govern

Classify the severity, assign decision rights, speak with one voice.

5 Recover

Return to service only when every restoration criterion is met: control, not speed.

6 Learn

An AI TRACE review within 5 to 10 business days; controls, model and board report updated.

Severity is classified on V-SEV, from V1 Irregularity to V5 Systemic Trust Event, where confidence in AI governance itself is at risk. The rule is simple: when in doubt, classify up. In critical national infrastructure the AI Office adds its own clock: exposure known within 4 hours, harm contained within 24, a decision on whether it is safe to continue within 72, and lessons in the controls within 14 days. The board sets these times.

For the board Rehearse an AI incident with the executive team this year, and set the board's notification threshold by severity.

11 · Where it fits

The model scales from one operator to a whole investment portfolio

Illustration · Power, rail, water and telecoms: essential services joined by one network



Critical national infrastructure

Operators of essential services carry duties under the Network and Information Systems Regulations 2018 and their sector regulator, with wider duties expected from the Cyber Security and Resilience Bill. Here the AI Office ties AI into the existing safety case, management of change and incident reporting, so the regulator sees one coherent picture rather than a separate AI story.

Regulated industry

Where named senior accountability already exists, as in financial services, the AI Office gives each accountable executive the evidence they need: what AI touches their area, who signed it and how it is performing.

Venture capital and private equity portfolios

A fund can run one AI Office across its holdings. Portfolio companies are too small to build this alone, and AI risk is correlated: many depend on the same few models and clouds, so one supplier failure can hit several holdings at once. The fund holds a hub; each company names an AI lead; each company's board stays accountable.

Proportion across a portfolio

Tier	The company	What applies
1	AI is the product, or the sector is regulated or safety-critical	The full approach, including security assurance and the sign-off chain

Tier	The company	What applies
2	AI runs core operations or customer decisions	Register, named ownership, capability score, value gates, security testing on material uses
3	AI is used as a productivity tool	Register, approved tools, staff rules and an annual review

The payoff comes at exit. A buyer pays more for AI that is owned, tested and documented, and discounts AI it cannot verify.

For the board Decide the level of governance each business or holding needs. Not every one needs the full model.

12 · Getting started

A working AI Office in 90 days, proven within a year

Velinor stands up the core functions in the first 90 days, runs them alongside the client's people, and hands over a working office with its own evidence, not a report about one. The full plan runs over twelve months in four phases. Each phase ends with evidence, not a status report, and capital follows the gates.

Figure 4. Twelve months in four phases, drawn to scale

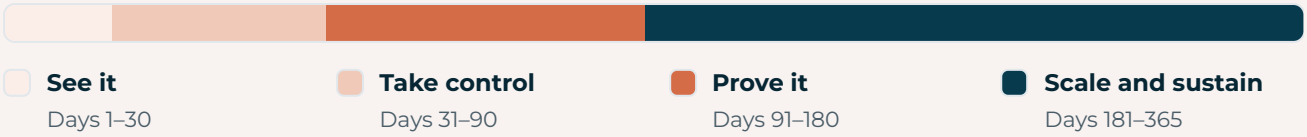

Four phases

Phase	Days	Ends with
See it	1 to 30	A complete AI register; named outcomes with measured baselines; a first likelihood of loss; a VTF baseline; the board's G0 decision
Take control	31 to 90	Risk appetite in numbers; three names against every registered use; release standards published; top risks restated in numbers; directors briefed
Prove it	91 to 180	One outcome where the work has changed and the measure has moved; live tests passed; first forecasts scored; the G1 decision

Phase	Days	Ends with
Scale and sustain	181 to 365	Remaining outcomes scaled or stopped; independent review of controls; first AI section in the ESG report; G2 and G3 decisions

The first ten days

1. The chief executive announces the AI Office and names the Senior Responsible Owner.
2. The chair names a non-executive lead for AI.
3. The finance director confirms what the AI budget covers and what is already committed.
4. The Head of AI sends the register request to every business unit.
5. A calibration session is booked for the people who will give estimates.
6. A board date is set for gate G0.

For the board Set the date of the first value gate decision now. It fixes the pace of everything else.

13 · Board questions

Eight questions that expose a weak answer

A board does not need to understand the models. It needs to ask questions that expose weak answers.

Questions and red flags

Question	Red flag
What AI do we run, including tools nobody sanctioned?	No register, or one that covers only funded projects
Who is accountable, by name, for each AI-supported decision?	A committee named as owner
What likelihood, period and loss range sits behind "high"?	A colour with no number
How likely is our AI investment to lose money?	One benefit figure, no range
What is the track record of the people who gave these estimates?	No forecast has ever been scored

Question	Red flag
Which use case has actually changed how the work is done?	Usage figures offered instead of results
What happens to the service if this AI is switched off today?	No tested fallback
What evidence, beyond our own opinion, shows our governance improves decisions?	“We have a framework” with no measured result

For the board Use these questions at the next board meeting where AI spend or AI risk is on the agenda.

14 · Success

What good looks like at twelve months

- The board can state the likelihood that its AI investment returns more than it cost, and it is inside appetite.
- Every AI solution in service has a named owner, a tested fallback and a recorded release decision.
- At least one outcome has moved its measure, tracked against a forecast made in advance.
- Risks are stated as likelihoods and loss ranges, and the forecasts behind them have been scored.
- The people using AI are trained for the task, can override it and trust it for good reason.
- Suppliers are mapped and under contract terms that give the organisation control.
- The regulator can be shown, on request, who decided what and on what evidence.

That is reliable, resilient and responsible AI. It is what gives leaders clarity, control and confidence.

How Velinor can help

Start with the picture, then build the office around it

Velinor helps boards and executive teams govern AI so that it is reliable, resilient and responsible. We work alongside your people and leave behind a working capability with its own evidence, not a report about one.

AI Trust Audit

A fast, evidence-based view of where AI is used, who owns it and where the organisation is exposed, scored against the Velinor Trust Framework.

AI Office set-up

The core functions stood up in 90 days, run alongside your people, then handed over with the register, standards and value gates in place.

Fractional CAIO

Senior AI leadership on a part-time basis: a 180-day programme that takes the board from first picture to governed, measured AI.

About the author



James A Lang

Founder, Velinor · Fractional Chief AI Officer

James founded Velinor to help leaders use AI they can trust. He works with boards in critical national infrastructure, defence and regulated industry, and sits on the advisory board of D3IP, a network of around 250 defence companies.

Notes and sources

The financial figures in section 07 are illustrative. They come from invented input ranges, 200,000 simulations over seven years at a 3.5% discount rate, and describe no client. A real model uses the client's own calibrated ranges.

Douglas W. Hubbard, *The Failure of Risk Management* (Wiley, 2009). The method is applied in Velinor's words; its application to AI is Velinor's extension, not the author's claim.

NCSC and partners, *Guidelines for secure AI system development*, version 1.0, November 2023: the source of the four lifecycle stages.

Onora O'Neill, *A Question of Trust*, BBC Reith Lectures, 2002: the source of honesty, reliability and competence as the parts of trustworthiness.

"AI Office", oversight modes, value gates and Critical Information Requirements are common terms used here in a Velinor way. "AI Office" means the organisation's own team, not the European Commission's AI Office.

Important notice

This paper provides general information about the governance of artificial intelligence. It is not legal, regulatory, financial or investment advice, and it does not take account of any organisation's particular circumstances. Seek advice specific to your situation before acting on it.

The financial figures in section 07 come from an illustrative model built on invented input ranges. They describe no client or real organisation and are not a forecast. Laws, regulations and standards are named as summaries only; confirm the current text with your advisers at the date of use.

© Velinor AI Ltd, 2026. All rights reserved. R³AI, the Velinor Trust Framework, AIBlindspot, V-AIM, V-SEV and AI TRACE are Velinor frameworks.

Annex · The fourteen principles

What the AI Office delivers for the board: fourteen principles in four groups

The board owns the principles and sets the appetite: how much risk it will accept in return for the value it wants. The AI Office ensures, on the board's behalf, that every principle is executed within that appetite, and tells the board when one is not. Each principle comes with a test the board can apply.

Value We know what AI is worth, and we get it

No.	Principle	The board's test
1	Protect the investment and realise its value	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

No.	Principle	The board's test
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

No.	Principle	The board's test
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

No.	Principle	The board's test
11	Keep AI human-centred and safe	Who is the person in control of each system, 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?

Annex · The fourteen principles

One named lead answers for each principle, and the evidence reaches the board along one strand

Who leads each principle

AI Office function	Led by	Leads principles
Leadership of the AI Office	Senior Responsible Owner	4
Finance office	Finance director, with the contracting officer	1, 2
Assurance and ethics	Chief risk officer	3, 7
Operational delivery and watch	Head of engineering, with the operations director	6
Transformation and architecture	Transformation and communications officer	5
Security and resilience	Head of security	8, 9, 10
Safety	Head of safety	11
Data	Head of data	12, 13
Learning, training and people	Head of learning, training and people	14
Digital (proposed)	Head of digital	None yet

Two things stay with the board and cannot be passed down: the power to stop a system, and the duty to set the appetite.

The golden strand

Every piece of assurance work leaves a product. Each product is owned by one AI Office function, is the evidence for one or two principles, and becomes a line in the board's quarterly report. If a piece of work cannot be traced to a principle, ask why it is being done. If a principle has no evidence, the board is governing on trust.

1
Work
Govern, measure, assure, test and watch.

2
Product
A register, a map, a record or a brief.

3
AI Office function
A named lead owns it and checks it against the appetite.

4
Principle
The product is the evidence for a principle.

5
Board report
Status, evidence and the decision asked of the board.

For the board

Adopt the fourteen principles by resolution, set the appetite in numbers, and ask for a report against each principle every quarter.

The principles are set out in full, with the appetite, the information and the people a board needs, in the companion Velinor white paper Fourteen principles for governing AI, November 2026. Twelve of the principles take the Institute of Directors business paper AI Governance in the Boardroom as their reference point.



NEXT STEP

Set up your AI Office

Start with one question: what AI do we run, and who answers for it?
Let us help you set up your AI Office.

Book a free 30-minute scoping call

Scan the code, or write to james@velinor.io

James A Lang
Founder and CEO, Velinor
james@velinor.io · velinor.io

© Velinor 2026. This paper is general guidance and is not legal advice.