


A blank cell is not a gap in the drawing.

Layer-cake diagrams let you finish a tier and climb past it. Three things never finish: **human capability, governance, supply chain**. They run through every layer, and all three are required at each one. Read a row to see what a layer needs on every column. Read a column to see how one concern changes shape as you climb.

	HUMAN CAPABILITY	GOVERNANCE	SUPPLY CHAIN
Runtime monitoring & response	Tells drift from noise and actually overrides the model. Resists automation bias.	Live oversight and incident ownership, each owned by a name. Anchor signal: AI Oversight Override Rate	Vendor ships a new version silently. What can you see, and what can you stop?
Application the shipped surface	Scopes the use case honestly. Knows the limits of guardrails and the misuse patterns.	Use-case approval and the execution boundary — enforced structurally, owned, change-controlled. Permission is not judgment	Plugins, APIs, embedded agents. Can you constrain a copilot you did not build?
Model assurance the release gate	Reads an evaluation and designs a red team. Tells a real evaluation from a green dashboard.	Documented evaluation as the control, and a named go/no-go authority empowered to say no .	What you verify versus what you inherit. Can you even red-team a closed model?
Data build order starts here	Judges provenance, rights and representativeness. Knows fit-for-purpose from merely large.	Classification, lawful-basis sign-off, rights and retention owned by a name.	Are the rights to use this data for training contractually established, or assumed?

THE L — BASELINE LITERACY UNDERNEATH, ROLE-SPECIFIC SKILLS ALONGSIDE

The L has two arms. **Baseline AI literacy** is the organisation-wide floor everyone shares. **Role-specific AI skills** are that floor applied as the competence each layer demands — the twelve cells above are that arm written out. Neither is a foundation tier you complete and leave below. Without both, all three columns fail quietly: a governance owner who cannot read a failure, a reviewer who cannot judge an evaluation, a vendor assessor who cannot read a model card.



WALK ANY AI INITIATIVE THROUGH THE GRID AND ASK THREE THINGS PER CELL

01

Is it present **here**? Not “we did that at the data phase.”

02

Who is the **named person** who owns it at this layer?

03

What **cadence** tells the board it still holds?

A cell that answers “we handled it at the bottom” is a column that stopped running above that point.

Source: *A Practitioner’s Guide to AI Governance* — a framework by Fredrik Lindstrom. Columns are human capability, governance and supply chain; all three required. Compliance is not a column: it is what the governance column produces once every cell is owned and reported. Oversight is not a fifth layer, it is the governance column. **Date discipline:** EU AI Act Art 4 (AI literacy) in force 2 February 2025. Do not anchor on 2 August 2026 and never describe the Act as fully operational. Exact NIST subcategory and ISO Annex A codes live in the locked verified crosswalk. Status labels and column assignments above are this sheet’s reading of the cited source, not the source’s own 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.

ARTICLES & BOARD ADVISORY

fredriklindstrom.info

Long-form governance analysis and toolkits

LINKEDIN

in/fredriklindstrom

Daily commentary and article launches

YOUTUBE

@Promise-and-Risk-AI

AI education without the hype or fear

COMPANION SITE

promiseandrisk.ai

Episode notes, sources and explainers

NEWSLETTER

The Governance Memo

Monthly, board-facing. Subscribe at fredriklindstrom.info