

AI Governance Questions for Public Leaders

A four-page board briefing — the same checklist, condensed for a leadership-team or board agenda.

THE THESIS

Better questions, before better tools. AI deployment in a public institution is not primarily a technology choice. It is a governance choice — about responsibility, procurement power, data legitimacy, auditability, resilience and democratic control.

AI decisions are institutional decisions. Conformity assessments, model cards and vendor assurances do not transfer accountability away from the institution.

Procurement is governance. The contract is where most sovereignty, accountability and exit decisions are actually made — or quietly given away.

TOOL ADOPTION

What people usually ask

Can staff use this? Which tasks can be automated? How do we save time? What licence tier do we need?

GOVERNANCE

What leaders need to ask

Who is accountable? What data is used, on what basis? What must remain human-controlled? What dependencies are created? How do we audit, explain and exit?

Five decisions leadership cannot delegate.

Before any AI deployment is signed off, the leadership team should agree on these five — by name, in writing.

1

What is the institutional purpose?

A single sentence, written before the technology choice, that the head of the institution can defend in a parliamentary or board hearing. "Efficiency" alone is rarely sufficient.

2

Who is accountable?

A named individual at each accountability layer — not a function, not a committee, not a contract. If the answer requires a diagram, the diagram is the problem.

3

What must remain human-controlled?

An explicit list of decisions, populations, edge cases or workflows that the AI may inform but never determine. Average performance is not the right measure for decisions that touch rights, benefits or safety.

4

What dependencies are acceptable?

A clear-eyed list — vendor, model, cloud, jurisdiction, infrastructure layer, sub-processor — and an explicit decision to accept each one. Plus a documented exit position for the dependencies the institution cannot afford.

5

How will the institution verify, explain and exit?

Three named mechanisms — for legitimacy, for quality, for resilience — each on a defined cadence, each with an owner, each producing evidence the institution can show on request.

Nine areas. One page.

A leadership team can score each of the nine governance areas green, yellow or red — with one line of evidence beside each. The map matters more than the score.

<p>● Green</p> <p>Clear enough to proceed. Owners named, evidence available.</p>	<p>● Yellow</p> <p>Important gaps. Proceed with documented mitigation and a review date.</p>	<p>● Red</p> <p>Do not proceed before leadership review and corrective action.</p>
---	---	---

For each colour, write the evidence behind it. A status without evidence is only a feeling.

01	Purpose & mandate	A single sentence the head of the institution can defend, written before the technology choice.
02	Responsibility	A named individual at each accountability layer — not a function, not a committee, not a contract.
03	Data & legality	Lawful is the floor. A use the institution would not publish is not yet legitimate.
04	Procurement	An exit clause the institution could realistically execute, and a documented annual portability test.
05	Risk & failure modes	A short list of named failure scenarios with detection mechanisms, owners and rehearsed responses.
06	Human oversight	Reviewers with the time, skill and authority to disagree — and an auditable record of when they did.
07	Transparency	A documented case where an explanation was produced on request, with the actual logs to support it.
08	Resilience	An annual continuity drill in which the AI component is removed from a workflow — successfully, or visibly not.
09	Public trust	A self-initiated public statement that survives the first round of citizen, media and parliamentary questions.

From briefing to working tool.

A 90-minute scan can be run with the standing leadership team and a printout of this briefing. The output is a one-page map, an owners list, and a calendared review date.

90-minute scan

0–10 min	Why AI governance is now a leadership issue, not a technology one.
10–25 min	Use-case and dependency map — what AI is already in the institution, where, and on whose responsibility.
25–55 min	The five leadership decisions, applied to the most consequential live deployment.
55–80 min	Green, yellow, red scoring across all nine areas — with one line of evidence per colour.
80–90 min	Owners and next steps. One named owner and one named action per yellow and red, plus a review date.

What the institution should have at the end

- A one-page governance map of the most consequential AI deployment.
- A named owner for each yellow and red area.
- A short list of clauses required in the next AI procurement.
- A scheduled review date — quarterly is usually right.

For workshops, board briefings and public-sector AI governance support

Available for leadership briefings, board sessions, public-sector workshops and policy sparring — in English, Finnish or Swedish. The full briefing (twenty-page edition) is available on request and from juhanaharju.com.

juhanaharju.com · juhana@potamoi.eu · Helsinki · Brussels