BRIEFING · APRIL 2026 · VERSION 1.1

AI Governance Questions for Public Leaders

A decision-maker checklist for public institutions, boards and leadership teams.

Juhana Harju

AI, digital resilience and sovereignty briefings for decision-makers.

juhanaharju.com

Helsinki · Brussels

EXECUTIVE SUMMARY

Better questions, before better tools.

Public institutions are now in their third year of practical experimentation with generative AI, and in the first year of mandatory regulatory work under the EU AI Act. The conversation has begun to shift — from whether to deploy AI to how to govern its deployment. The shift is overdue and incomplete.

Adoption is the easy part. Buying licences, training staff, running pilots — these are real work, but they do not constitute governance. Governance is the institutional discipline that lets a minister, a board, a CIO, or an audit body answer one question: **when an AI system shaped a decision in our pipeline, can we explain it, justify it, and take responsibility for it?**

This checklist is built for the leaders who have to answer that question. It is not a technical specification, and it is not a legal compliance form. It is a structured set of questions a leadership team should be able to walk through before a procurement is signed, a deployment is scaled, or an AI policy is published. Each section can be used standalone in a thirty-minute scan or as the basis of a half-day workshop.

Key takeaways

1. **AI decisions are institutional decisions.** The choice to deploy AI is a governance choice before it is a technology choice.
2. **Procurement is governance.** The contract is where most sovereignty, accountability and exit decisions are actually made — or quietly given away.
3. **Data quality and legitimacy matter as much as model quality.** A high-performing model on the wrong data is a credibility risk, not a productivity gain.
4. **Human responsibility cannot be outsourced.** Conformity assessments, model cards and vendor assurances do not transfer accountability away from the institution.
5. **Resilience and exit power matter.** A system you cannot leave is a system that owns part of your decisions.
6. **AI literacy must become organisational capability.** One-off training does not produce judgment under uncertainty.
7. **Public trust depends on explainability, auditability and accountability.** All three have to survive a parliamentary committee, a journalist, and a citizen complaint.

START HERE

Five decisions leadership cannot delegate.

If a leadership team takes nothing else from this briefing, it should agree on these five questions — by name, in writing — before the next AI deployment is signed off.

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; it conceals the trade-offs that should be visible.

2
Who is accountable?

A named individual at each accountability layer — not a function, not a committee, not a contract. When the audit committee or a parliamentary inquiry asks who owned the decision, the institution should answer in one sentence.

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 any dependency the institution cannot afford to be locked into.

5
How will the institution verify, explain and exit?

Three named mechanisms — for legitimacy (public statement, external review), for quality (evaluation suite, monitoring, incident review), for resilience and exit (continuity drill, fallback test, alternative-provider rehearsal). Each on a defined cadence, each with an owner, each producing evidence the institution can show.

SCOPE NOTE

How this briefing relates to compliance.

A leadership tool that sits above and before the legal, technical and procurement detail — not a substitute for any of them.

This briefing is a **leadership-governance tool**. It is not a legal compliance checklist, a regulatory opinion, or a substitute for the technical, legal and procurement review that any specific AI deployment requires.

The EU AI Act enters into application in phases. General-purpose AI obligations began to apply in August 2025; high-risk system rules begin in August 2026. Implementing acts, harmonised standards and the first conformity assessments are still being negotiated. Member-state implementing measures, sector-specific rules — for health, finance, social security, justice, energy, defence — and the data-protection, public-procurement and security-accreditation regimes each impose their own obligations. None of those is replaced or summarised here.

The intention of this checklist is to sit above and before that detailed work. It is designed to help leadership ask the right questions before the institution locks itself into tools, vendors, workflows or dependencies that the later legal and technical review can no longer easily change.

Used well, the briefing complements — and reduces the cost of — your DPIA, your AI Act conformity assessment, your procurement-law review and your sector-specific compliance work. It does not replace any of them.

> Compliance work asks "Are we allowed to do this?" Governance asks "Should we, and on what terms?" Both questions need an answer before deployment. This briefing addresses the second.

HOW TO USE THIS CHECKLIST

A working tool, not a compliance form.

The checklist is designed to be used in the room, by the people who actually decide. Print it, share it, copy the questions into your own deck. It is most useful when it surfaces disagreement.

Use it in

- Leadership team meetings reviewing an AI initiative.
- Board discussions on technology strategy or risk.
- Procurement preparation, before an RFP is drafted.
- AI policy drafting, internal or sector-level.
- Internal risk and controls reviews.
- Pre-project assessment and stage gates.

Suggested formats

- **30-minute scan.** A leadership team reads the questions, marks each section green, yellow or red, and identifies the one section that needs the most work.
- **90-minute discussion.** Two or three sections worked through in depth, with explicit owners and next steps assigned.
- **Half-day workshop.** All nine sections, with a use-case map and a written governance posture as the output.

Working principles

- Do not try to answer every question perfectly. Use them to find blind spots.
- Mark each area green, yellow or red. The map matters more than the score.
- Assign an owner and a next step for every yellow and every red.
- Revisit on a defined schedule. Governance is a posture, not an event.

THE CORE FRAME

AI is not just a tool issue.

Most public-sector AI investment so far has gone into tool training. The harder, more consequential work — institutional governance — is too often distributed between consultants, vendors and regulatory interpretation. Ultimate responsibility, however, remains with the institution itself, and so does the need to ask the institutional questions first.

Tool training teaches people to use AI. Governance helps an institution decide how AI should be used. Both are necessary; the first is widely funded and the second is widely under-funded. Closing that gap is what this checklist is for.

> The strategic question is not only "What can AI do?" but "Who is responsible, what changes in the process, what dependencies are created, and how do we know the outcome is legitimate?"

Most institutions can answer the first question (what AI can do) within an afternoon. Few can answer all four of the second set within a quarter. The gap between those two timelines is where the governance work happens. Everything below this page is structured to close it.

TOOL ADOPTION

What people usually ask

- Can staff use ChatGPT or Copilot?
- What prompts work best?
- How do we save time?
- Which tasks can be automated?
- What licence tier do we need?

GOVERNANCE

What leaders need to ask

- Who is accountable for AI-assisted decisions?
- What data is being used, and on what legal basis?
- What must remain human-controlled?
- What vendor and infrastructure dependencies are created?
- How do we audit, explain and, if needed, exit?

01 CHECKLIST

Purpose and mandate.

Force clarity on why AI is being used — before any conversation about which model, which vendor, or which budget line.

A surprising number of AI deployments in the public sector cannot articulate, beyond generalities, what public purpose the deployment serves. "Efficiency" is rarely sufficient as a stated purpose; it conceals decisions that should be visible — about who benefits, what is being de-prioritised, and what trade-offs are being made. The first discipline of governance is forcing that clarity into the open, in writing, before the technology choice begins.

- ☐ What public purpose does this AI use serve?
- ☐ Is the purpose efficiency, quality, accessibility, risk reduction, better service — or something else?
- ☐ Is the use case aligned with the organisation's legal mandate?
- ☐ Is this a real organisational need, or technology-driven experimentation that has found a sponsor?
- ☐ Who has the authority to approve this use, and is that authority being exercised explicitly?
- ☐ What would be the cost of not using AI here?
- ☐ What would be the cost of using it poorly?
- ☐ How will success be measured, and by whom?
- ☐ What must not be compromised even if efficiency improves?

> THE TEST FOR THIS AREA
>
> A single sentence, written before the technology choice, that the head of the institution can defend in a parliamentary or board hearing — and that survives the first round of difficult follow-up questions.

02 CHECKLIST

Responsibility and accountability.

Make responsibility explicit, named, and exercisable — not implied, distributed or contractually deflected.

A common failure mode in public-sector AI deployment is not technical. It is the quiet diffusion of responsibility across legal, IT, procurement, data protection, the vendor, and the frontline user — until no individual or function can be said to own the AI-assisted decision. When something goes wrong, that diffusion becomes visible as a vacuum. Naming responsibility upfront is the only way to prevent it.

- ☐ Who owns the AI use case end-to-end?
- ☐ Who is accountable for the final decision or output?
- ☐ Who is responsible for monitoring quality on an ongoing basis?
- ☐ Who can stop the system or process if something goes wrong, and how quickly?
- ☐ Where does human judgment enter the process, and is that point load-bearing or theatrical?
- ☐ Is responsibility clear between leadership, legal, IT, procurement, data protection and frontline users?
- ☐ Are external vendors making decisions that should remain with the institution?
- ☐ Can accountability be explained — without footnotes — to citizens, auditors, media and oversight bodies?

> THE TEST FOR THIS AREA
>
> A named individual at each accountability layer — not "the team," not "the vendor," not "the system." If the answer to "who is responsible?" requires a diagram, the diagram is the problem.

03 CHECKLIST

Data, legality and legitimacy.

Connect data governance to public trust. Lawful is not the same as legitimate, and both are necessary.

Data is where AI deployments most often fail their public-trust test. A use that is lawful under GDPR can still be illegitimate in the eyes of the people whose data feeds the system. The two tests are different, and an AI deployment in the public sector has to pass both. The questions in this section assume the legal basis is the floor, not the ceiling.

- ☐ What data does the system use, and from where?
- ☐ Is the data lawful, relevant and necessary for the purpose declared in Section 01?
- ☐ Is sensitive or confidential data involved?
- ☐ Is personal data involved, and has a data protection impact assessment been completed?
- ☐ Is the data representative enough for the intended use, and across which subgroups?
- ☐ Are there known biases or gaps that change who the system serves well or poorly?
- ☐ Can citizens understand — in plain language — how their data is being used?
- ☐ Is there both a legal basis and a legitimacy basis?
- ☐ Would the use feel acceptable to the public if it were disclosed in tomorrow's headline?

> THE TEST FOR THIS AREA
>
> A short, public-facing description of what data the system uses and why, that a journalist could quote without misrepresenting the institution. If you would not publish that paragraph, the underlying use is not yet ready.

04 CHECKLIST

Procurement and vendor dependency.

Procurement is the single largest unforced governance decision in most institutions. The contract is where sovereignty is either kept or quietly traded away.

Most AI governance lives or dies in the procurement clause. A well-meaning policy is undone by a contract that locks the institution into a particular model, prevents auditing, makes data exit prohibitively expensive, or routes decisions through subcontractors the buyer never met. Procurement should be treated as the most important governance instrument the institution owns — because, in operational terms, it usually is.

- ☐ Are we buying a tool, a service, a model, a platform or an ongoing dependency?
- ☐ Where is data processed and stored, and under whose jurisdiction?
- ☐ Is data residency enough, or do we need stronger control — operational, key-management, jurisdictional?
- ☐ Can we audit the system, and what does the audit clause actually allow?
- ☐ Can we export our data and our workflows, or only the data?
- ☐ Can we change provider without unacceptable disruption — and have we tested it?
- ☐ Are model changes, updates and silent capability shifts contractually addressed?
- ☐ What happens if the provider changes pricing, terms, ownership or availability?
- ☐ Do we understand subcontractors and infrastructure dependencies all the way down?
- ☐ Are there lock-in risks that will not appear in the first invoice but will appear in the third?
- ☐ What clauses must be required in the contract before signing?

> THE TEST FOR THIS AREA
>
> An exit clause that the institution could realistically execute within a defined period, and a documented annual portability test. If exit is theoretical, the dependency is real.

05 CHECKLIST

Risk, safety and failure modes.

Move beyond abstract risk talk. Specify what can go wrong, who could be harmed, and what the institution will do when it does.

"AI risk" as a category is too broad to govern. What can be governed is a specific list of failure modes for a specific deployment, with named owners, defined detection mechanisms and pre-agreed escalation paths. The questions below force that specificity. A risk register that does not contain at least three concrete failure scenarios for the system in front of the institution is not yet a risk register in any operational sense.

- ☐ What can go wrong with this system, in concrete terms?
- ☐ Who could be harmed — directly, indirectly, individually, at scale?
- ☐ What are the most likely failure modes? What are the worst credible failure modes?
- ☐ Could errors affect rights, services, benefits, safety, or institutional trust?
- ☐ Can false outputs be detected — automatically, by a reviewer, by the affected party?
- ☐ Is there a fallback process, and has it been used recently enough to be functional?
- ☐ What happens during an outage or degraded performance?
- ☐ How are incidents reported, and to whom, and within what time?
- ☐ Who reviews incidents, decides what changed, and authorises the next iteration?
- ☐ Is there a clear escalation path — internal, regulatory, public-facing — and does it work on a Friday evening?

> THE TEST FOR THIS AREA
>
> A short list of named failure scenarios with detection mechanisms, owners, and rehearsed responses. If the list is generic, the risk work is not yet at operational depth.

06 CHECKLIST

Human oversight and organisational capability.

Avoid fake oversight. A reviewer who does not have the time, skill or authority to disagree with the system is not oversight — it is paperwork.

The phrase "human in the loop" appears in many more AI policies than deployment realities. A human reviewer who sees three hundred outputs an hour will rubber-stamp them; a reviewer who lacks the authority to overrule the system will defer to it; a reviewer who has not been trained on its failure modes will not see them. Real oversight requires time, skill, authority and continuous learning — and the absence of any one of those four converts oversight into theatre.

- ☐ What does human oversight mean in this specific use case — concretely, not in policy language?
- ☐ Does the human reviewer have enough time, skill and authority to disagree with the system?
- ☐ Are people genuinely able to challenge the AI output, including upward?
- ☐ Are users trained to detect errors, hallucinations and biased outputs?
- ☐ Are staff over-trusting the system, and how would the institution notice?
- ☐ Are staff under-using the system because of fear or uncertainty, and is that being addressed?
- ☐ What new skills are needed, and at which levels of the organisation?
- ☐ Is competence covered by one-off training or by continuous learning?
- ☐ Who maintains organisational AI competence as a function, not as a project?

> THE TEST FOR THIS AREA
>
> An auditable record of cases where a reviewer overruled the system — and a senior leader who can explain how the institution learns from those cases.

07 CHECKLIST

Transparency, auditability and explainability.

Connect technical systems to democratic accountability. The audit trail is where institutional AI either becomes legible — or quietly does not.

Public institutions are not private companies. The standards for transparency, auditability and explainability are higher because the legitimacy at stake is democratic, not commercial. The questions below are not about generating user-friendly explanations of model internals; they are about whether the institution can reconstruct, defend and disclose its own use of AI when asked. Many institutions today cannot, in any robust way. The fix is mostly procedural, not technical.

- ☐ Can we explain the AI use to citizens in plain language?
- ☐ Can we explain it to internal leadership, in operational terms?
- ☐ Can we explain it to an auditor or regulator, in evidentiary terms?
- ☐ Are decisions or recommendations logged — with prompt, model version, output and reviewer recorded together?
- ☐ Is there an audit trail that survives a personnel change and a cloud migration?
- ☐ Can we reconstruct, after the fact, why a particular decision was made?
- ☐ What information should be disclosed publicly, and in what cadence?
- ☐ What information must remain protected, and on what defensible basis?
- ☐ Is the system understandable enough for its level of impact?
- ☐ Are appeal and correction mechanisms clear to the person on the other end of the decision?

> THE TEST FOR THIS AREA
>
> A documented case where a citizen, journalist or auditor asked for an explanation, and the institution produced one within a defensible time, with the actual logs to support it.

08 CHECKLIST

Resilience and continuity.

Tie AI governance to digital resilience. A function that cannot be performed without AI is a function the institution no longer fully owns.

Every AI deployment quietly inherits the resilience profile of the cloud it runs on, the model it depends on, the vendor that maintains it, and the geopolitical relationship that allows them all to be procured. That inheritance is rarely made visible inside the institution. Resilience and continuity questions force it onto the table — before the next regional outage, model deprecation or contract dispute does it for them.

- ☐ What processes become dependent on AI — and how dependent?
- ☐ What happens if the AI tool is unavailable for an hour, a day, a week?
- ☐ What happens if cloud access is disrupted, regionally or globally?
- ☐ What happens if model quality changes suddenly between versions?
- ☐ Do we have a manual fallback, and has it been tested in the last twelve months?
- ☐ Do we have alternative providers we could realistically switch to?
- ☐ Do we know our critical dependencies all the way down to the chip layer?
- ☐ Are key workflows documented outside the tool, in a form a new staff member can follow?
- ☐ Can we continue operating during vendor, cyber, legal or geopolitical disruption?
- ☐ Is AI part of our continuity planning, or has it been silently exempted?

> THE TEST FOR THIS AREA
>
> An annual, documented continuity drill in which the AI component is removed from a workflow, and the institution operates without it for a defined period — successfully, or visibly not.

09 CHECKLIST

Public trust and democratic control.

Elevate the discussion beyond compliance. The point of public-sector AI governance is not to pass the audit — it is to keep the institution legible to the people it serves.

Compliance is necessary and insufficient. An institution that meets every regulatory requirement can still erode public trust, exclude vulnerable groups, or quietly substitute administrative judgment with model output. The questions in this final section reach past the compliance frame and into the democratic one. They are uncomfortable, which is the point.

- ☐ Would citizens find this use of AI legitimate, if they understood it?
- ☐ Would the institution be comfortable explaining the use publicly, in detail, on its own initiative?
- ☐ Does the use strengthen or weaken trust in the institution?
- ☐ Could the system create hidden policy choices — by selection, ranking or thresholding — that should be visible?
- ☐ Are elected decision-makers and oversight bodies able to understand the implications?
- ☐ Is meaningful democratic oversight possible, given the system's architecture?
- ☐ Are vulnerable groups affected, and have they been consulted?
- ☐ Are there equality, access or language issues the deployment creates or worsens?
- ☐ Is the institution still visibly responsible — or has the AI component become a polite way of saying "the system decided"?
- ☐ Does the system support human dignity and fairness, or merely fail to violate them?

> THE TEST FOR THIS AREA
>
> A public statement of how the institution uses AI — written on the institution's own initiative, not in response to a complaint — that survives the first round of citizen, media and parliamentary questions.

PRACTICAL EXAMPLE

AI-assisted permit handling.

A deliberately ordinary public-sector scenario, walked briefly through the nine governance areas. The point is to show what "good enough to proceed" looks like in practice.

A municipal building-control office receives roughly 6,000 permit applications per year. Caseworkers spend a substantial share of their time on intake — checking that an application is complete, classifying it under the correct permit category, and identifying the relevant prior cases and ordinances. The institution is considering an AI-assisted intake tool that triages incoming applications, classifies them, and drafts a preliminary recommendation memo for the caseworker to review before the formal decision.

The nine governance areas, applied:

01 · Purpose and mandate. Stated purpose: reduce intake processing time and improve consistency of classification. The institution writes the purpose down — and explicitly declines to extend the system to final permit decisions, which remain with the caseworker.

02 · Responsibility. The head of building control is accountable for the deployment. Each AI-assisted recommendation is signed off, in name, by the responsible caseworker before it leaves the office.

03 · Data and legality. Inputs: incoming applications and the office's own historical decisions. No special-category personal data. A DPIA is completed; the legal basis is documented; applicants are informed in plain language that AI tools are used in intake.

04 · Procurement. Vendor selected through an open RFP that includes named clauses for audit, data export, model-version logging, exit notice and a cap on switching costs. Hosting in an EU data centre under EU jurisdiction; key management in a customer-controlled HSM.

05 · Risk and failure modes. Identified failure modes: misclassification (sending an environmentally sensitive case down the wrong path), hallucinated regulatory references, silent model drift. Each has a named detection mechanism and a fallback to manual intake within 24 hours.

06 · Human oversight. Caseworkers receive structured training on the system's known failure patterns, are explicitly authorised to overrule any AI-suggested classification, and a sample of overruled cases is reviewed quarterly to detect over- or under-trust.

07 · Transparency. Each AI-assisted draft is logged with input, model version, recommendation and the caseworker's final decision. The institution publishes a short public statement of how AI is used in intake.

08 · Resilience. Manual intake is documented as a fallback procedure and tested annually. The vendor's regional outages are mapped; an alternative-vendor exit is rehearsed.

09 · Public trust. A public-facing description of how AI is used, who is accountable, and how applicants can request a human review is published on the municipality's website. Annual reporting includes the AI-assisted share of intake, the override rate, and material incidents.

SAMPLE SCORING

Green across **01–04** and **07–09**; yellow on **05** (drift detection still maturing) and **06** (training scheduled but not yet delivered to all staff). The institution proceeds with a defined six-month review point and named owners for the two yellows.

SCORING TOOL

Green, yellow, red.

A simple way to convert nine sections of questions into a one-page leadership view — and a list of next steps.

The point of the scoring is not the score. It is to surface where the leadership team agrees, where it disagrees, and which areas need work before the institution proceeds. Use the categories below; do not over-engineer them.

Green

Clear enough to proceed or pilot carefully. Owners named, evidence available, no unresolved disagreement.

Yellow

Important gaps remain. Proceed only with documented mitigation, named owner, and a defined review date.

Red

Do not proceed before leadership review and corrective action. Escalate explicitly; document the decision either way.

> **For each score, write down the evidence behind the colour.** A status without evidence is only a feeling.

AREA	STATUS	OWNER	NEXT STEP
01 · Purpose and mandate	○ G ○ Y ○ R		
02 · Responsibility	○ G ○ Y ○ R		
03 · Data and legality	○ G ○ Y ○ R		
04 · Procurement	○ G ○ Y ○ R		
05 · Risk and failure modes	○ G ○ Y ○ R		
06 · Human oversight	○ G ○ Y ○ R		
07 · Transparency	○ G ○ Y ○ R		
08 · Resilience	○ G ○ Y ○ R		
09 · Public trust	○ G ○ Y ○ R		

SUGGESTED FORMATS

Leadership workshop agenda.

Two formats — a 90-minute working session and a half-day deep-dive — usable in ministerial, regulatory, agency and board settings.

90-minute version

0–10 min	Why AI governance is now a leadership issue, not a technology one.
10–25 min	Use-case and dependency mapping — what AI is already in the institution, where, and on whose responsibility.
25–45 min	Responsibility, data and procurement questions (Sections 02, 03, 04).
45–65 min	Risk, oversight and resilience (Sections 05, 06, 08).
65–80 min	Green, yellow, red scoring across all nine areas, with evidence noted for each colour.
80–90 min	Decisions, owners and next steps. One named owner and one named action per yellow and red.

Half-day version

Block 1	Strategic context — the EU AI Act, the open-weight floor, sovereign-cloud landscape, and what they mean for this institution.
Block 2	Organisational AI maturity — what is in production, in pilot, in shadow IT.
Block 3	Use-case selection and prioritisation — which deployments deserve governance attention first.
Block 4	Governance model — accountability, oversight, escalation, evaluation cadence.
Block 5	Procurement and vendor control — clauses, exit positions, dependency map.
Block 6	Risk and resilience — failure modes, continuity drills, escalation paths.
Block 7	Roadmap — twelve-month governance plan, with named owners and review dates.

"Good AI governance does not begin with fear or hype. It begins with better institutional questions."

Use this briefing to

— Run a 90-minute AI governance scan with your leadership team.
— Prepare an AI procurement or RFP before vendor engagement.
— Review an existing AI pilot before scaling it institution-wide.
— Build a 12-month institutional AI governance roadmap.
— Design a board or executive workshop on AI responsibility, resilience and public trust.

About the author

Juhana Harju is an independent voice on European technology policy — artificial intelligence, digital resilience, digital sovereignty and evidence-based policy. He is the sole author of Finland's 2026 ministerial report on new technologies and digital resilience, Head of Growth & Strategy at Solidate Oy, and a member of the advisory board of Statistics Finland. Earlier roles include founding CEO of NORDXE, a regulated Nordic digital-assets exchange, and CEO of Potamoi Group. He serves as secretary of the technology working group of the Social Democratic Party of Finland and co-founded tietopolitiikka.fi, a cross-partisan platform for evidence-based information policy.

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.

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