AI ADOPTION IN INSTITUTIONS

AI READINESS CHECKLIST

GUENIX

AI READINESS CHECKLIST

for Public Institutions, NGOs, and Mission-Driven Organisations

A governance prerequisite assessment for institutional leaders considering AI adoption.

INTRODUCTION — PURPOSE AND LIMITS OF THIS CHECKLIST

This checklist is a diagnostic instrument. It is designed for one purpose: to determine whether the minimum governance conditions for responsible AI deployment are currently in place at your institution.

It is not a guide to AI implementation. It does not explain how to correct the conditions it identifies. It does not evaluate the technical capabilities of any AI system. It does not assess vendor proposals.

It is addressed to executive directors, programme directors, governance officers, board members, and any institutional leader who bears fiduciary responsibility for the decisions their organisation makes and the populations it serves.

> **Who should complete this checklist**
> *This checklist should be completed before any AI procurement or deployment decision is made. It should be completed by the individual who holds governance responsibility for the initiative, not the individual who is technically implementing it. Answer based on current operational conditions, not on policies that exist on paper but are not followed in practice.*

The checklist contains 25 questions organised across four governance dimensions. Each question requires a Yes, No, or Not Sure response.

Questions marked ⚠ are critical. A single 'No' response on a critical question means the institution is not ready for deployment, regardless of the answers to all other questions.

A 'Not Sure' response on any question indicates that the governance condition has not been verified. An unverified condition cannot be assumed to be met.

A note on honest assessment

The most common failure in readiness assessments is describing institutional conditions as they are intended to be rather than as they are in practice. A data governance policy that exists on paper but is not followed in practice does not meet the standard. An override protocol that has been drafted but never communicated to staff does not meet the standard. Answer what is, not what should be.

DIMENSION 1 — DATA INTEGRITY AND COHERENCE

AI systems operate exclusively on data that is available to them. They cannot detect missing information, suspect incoherence, or seek context outside their data environment. Fragmented, inconsistent, or incomplete data does not produce imprecise outputs. It produces precise outputs based on a distorted picture of reality. The following questions assess whether your institution's data architecture can support responsible AI deployment.

Semantic integrity — definitions and meaning

Semantic fragmentation occurs when the same term is used to mean different things across departments. A predictive model trained on data from multiple sources will calculate correlations between concepts that do not actually correspond, producing precise results based on a linguistic misunderstanding.

1.1	**⚠ Has the institution identified all core terms such as beneficiary, risk, household, eligibility, or incident that the AI system will use, and verified that these terms are defined consistently across all departments whose data the system will access?**	☐ Yes	☐ No	☐ Not sure
1.2	Does a maintained Data Dictionary exist that records the agreed definitions of these terms, and is it used in practice rather than filed as a reference document?	☐ Yes	☐ No	☐ Not sure
1.3	**⚠ Is there a named authority, a Data Governance Committee, Chief Data Officer, or equivalent with the mandate to resolve definitional conflicts between departments when agreement cannot be reached voluntarily?**	☐ Yes	☐ No	☐ Not sure

Temporal integrity — timestamps and causality

Temporal fragmentation occurs when data entry reflects when staff had capacity to record events rather than when events occurred. A model trained on this data will learn patterns from administrative processing rhythms rather than operational reality and will encode false causal relationships that cannot be corrected by cleaning the data.

1.4	**⚠ Has the institution measured the average lag between events occurring and being recorded in the systems the AI**	☐ Yes	☐ No	☐ Not sure

	model will use and assessed whether this lag is sufficient to distort causal inference?			

1.5	Do the institution's data collection processes capture both the date an event occurred and the date it was recorded as separate fields rather than recording only the entry date?	☐ Yes	☐ No	☐ Not sure

Structural integrity — connectivity and completeness

Structural fragmentation occurs when records that logically belong together are physically separated across different systems with no common unique identifier. A model operating on structurally fragmented data will make confident assessments of situations it has never fully seen.

1.6	**⚠ Does a shared unique identifier exist across the systems whose data the AI model will need to access, allowing records about the same individual, household, or case to be connected without manual reconciliation?**	☐ Yes	☐ No	☐ Not sure
1.7	**⚠ Has the institution mapped every system that holds records relevant to the AI system's decisions, and assessed whether each system can be connected technically, legally, and institutionally?**	☐ Yes	☐ No	☐ Not sure
1.8	Where full connectivity between systems is not achievable, has the institution formally documented which information the model will not be able to access and the implications for the accuracy and fairness of its outputs?	☐ Yes	☐ No	☐ Not sure

DIMENSION 2 — GOVERNANCE AND DECISION RIGHTS

The most consequential governance failures in AI adoption are not technical. They are structural: the absence of defined authority, the silence on liability, and the paralysis that results when staff face accountability for decisions they did not make but cannot contest. The following questions assess whether your institution has resolved the governance questions that must be answered before a machine is permitted to influence institutional decisions.

The governance gap — authority and accountability

The governance gap is the structural distance between acquiring an AI capability and possessing the institutional readiness to manage it safely. It manifests as absent liability policies, undefined decision rights, and inadequate data governance. Its presence is often invisible until a failure makes it visible at speed and scale.

2.1	**⚠ Does the institution have a documented policy specifying who bears institutional liability for decisions made on the basis of AI system outputs, including decisions that follow the system's recommendation without human modification?**	☐ Yes	☐ No	☐ Not sure

2.2	**⚠ Has the institution identified the specific decisions the AI system will influence and documented which institutional role bears responsibility for each category of decision?**	☐ Yes	☐ No	☐ Not sure

2.3	Does the institution have a Data Governance Committee or equivalent authority with a defined mandate that covers AI governance not merely data storage and access?	☐ Yes	☐ No	☐ Not sure

Human override — authority, protection, and protocol

The liability freeze occurs when staff can technically override an AI recommendation but face personal accountability for incorrect overrides while facing no accountability for incorrect deference. This asymmetry produces systematic deference to algorithmic outputs regardless of their quality because following the machine is always defensible, and contradicting it never is.

2.4	**⚠ Is there a formal protocol specifying which roles hold the authority to override an AI recommendation, under what conditions an override is permissible, and what documentation is required?**	☐ Yes	☐ No	☐ Not sure

2.5	**⚠ Is an officer who overrides an AI recommendation in accordance with the protocol formally protected from**	☐ Yes	☐ No	☐ Not sure

	adverse audit or accountability consequences for that decision?			

2.6	**⚠ Does a documented escalation pathway exist for cases that fall outside the model's parameters, produce outputs that contradict professional judgment, or require a decision the system is not designed to make?**	☐ Yes	☐ No	☐ Not sure

Grievance and affected populations

When an AI system produces an incorrect decision affecting a real person, that person must have a mechanism to contest it. The absence of a grievance mechanism does not prevent harm from occurring. It prevents the institution from detecting it.

2.7	**⚠ Does a grievance mechanism exist for individuals affected by an algorithmic decision allowing them to contest that decision and receive a documented human review?**	☐ Yes	☐ No	☐ Not sure
2.8	Has the institution assessed which populations are most likely to be incorrectly classified or excluded by the AI system and put in place monitoring to detect systematic exclusion patterns?	☐ Yes	☐ No	☐ Not sure

Procurement — contractual governance

Procurement is a governance decision. A contract that does not specify data provenance, inspection rights, liability allocation, and model update notification does not transfer governance risk to the vendor. It concentrates it silently on the institution.

2.9	**⚠ Does the vendor contract require the vendor to disclose which datasets trained the model, their temporal range, and any known limitations or biases identified during development?**	☐ Yes	☐ No	☐ Not sure
2.10	**⚠ Does the contract grant the institution the right to inspect the model's training data, architecture, and performance logs at any point during the contract period?**	☐ Yes	☐ No	☐ Not sure
2.11	Does the contract require the vendor to notify the institution before deploying any material change to the model ? Does the institution retain the right to refuse or delay that change?	☐ Yes	☐ No	☐ Not sure
2.12	Does the contract explicitly allocate liability for decisions made on the basis of model outputs, and specify where institutional data is stored and under which legal jurisdiction?	☐ Yes	☐ No	☐ Not sure

DIMENSION 3 — AUDITABILITY AND ACCOUNTABILITY

An institution that cannot explain how a decision was made, which version of a model was used, on the basis of which data, and under which governance conditions cannot be held accountable for it. Auditability is not a technical feature of the software. It is an institutional obligation that must be designed, maintained, and verified independently of the vendor.

Audit trail — retrievability and version control

An audit trail is not a log. It is a structured record that allows any specific decision to be reconstructed including the data inputs used, the model version active at the time, and the output produced. Without version control and data retention, reconstruction is impossible and accountability is fictional.

3.1	**⚠ Can the institution reconstruct, for any specific decision the AI system has produced, the exact data inputs used, the version of the model that was running, and the output it produced?**	☐ Yes	☐ No	☐ Not sure
3.2	**⚠ Are governance artefacts such as data dictionaries, decision rights protocols, override logs, bias assessments and model version records controlled, dated, attributed to a named authority, and retained in a format accessible to external auditors?**	☐ Yes	☐ No	☐ Not sure
3.3	Is a named role within the institution responsible for maintaining governance artefacts and conducting formal auditability reviews at defined intervals, with the mandate and resources to do so independently?	☐ Yes	☐ No	☐ Not sure

Explainability — decisions and affected individuals

The minimum standard of explainability is this: any decision the AI system produces must be explainable, in terms accessible to the individual affected, by a human officer of the institution. A system that cannot meet this standard cannot be governed by the institution that deploys it.

3.4	**⚠ Has the institution verified that its officers can explain, in plain language accessible to the individual affected, why the AI system produced a specific output, not merely what the output was?**	☐ Yes	☐ No	☐ Not sure
3.5	**⚠ Has a bias assessment of the model's training data and outputs been conducted and documented before deployment, including an assessment of differential accuracy across demographic groups?**	☐ Yes	☐ No	☐ Not sure

DIMENSION 4 — INSTITUTIONAL TRUST AND OPERATIONAL FIT

A system that staff do not trust will not be integrated. It will be installed. The distinction matters: installation is a technical event that makes a capability available. Integration is an institutional process that makes a capability genuinely used and relied upon. A system that staff bypass, work around, or comply with performatively has not been integrated. It has created a governance liability without delivering operational value.

Informal governance — workarounds and discretion

Every institution relies on informal practices to function. These include discretionary exceptions, verbal approvals, and policy workarounds that bridge the gap between documented rules and operational reality. AI systems eliminate these practices. Institutions that have not mapped them prior to deployment will discover them only after the system has exposed their absence at speed and at scale.

4.1	**⚠ Has the institution mapped the informal governance practices that currently sustain the process being automated, including undocumented workarounds, discretionary exceptions and verbal approvals?**	☐ Yes	☐ No	☐ Not sure

4.2	For each informal practice identified, has the institution made a deliberate decision to either formalise it, preserve it through human-in-the-loop design, or accept that its elimination, with documented acknowledgment of the consequences?	☐ Yes	☐ No	☐ Not sure

Staff readiness — confidence and clarity

Staff who do not understand what an AI system can and can not do, who do not know their override rights, and who are not protected from the consequences of exercising professional judgment will not engage with the system honestly. They will defer, bypass, or comply performatively. Each of these creates a distinct category of governance failure.

4.3	**⚠ Have the staff who will use or be affected by the AI system been informed of the conditions under which its outputs should be trusted, the conditions under which they should be questioned, and their right to escalate?**	☐ Yes	☐ No	☐ Not sure

4.4	Have the staff who hold override authority been explicitly informed of that authority, the conditions under which it applies, and the documentation required when it is exercised?	☐ Yes	☐ No	☐ Not sure

Pilot and deployment conditions

A successful pilot does not confirm deployment readiness. It confirms that the technology functions under controlled conditions. The relevant question is whether the institution, not the technology, is ready to absorb it under normal operational constraints.

4.5	Has the institution assessed whether the conditions under which the pilot was conducted, curated data, volunteer participants, and dedicated technical support, are realistically replicable in the full deployment environment?	☐ Yes	☐ No	☐ Not sure

INTERPRETING YOUR RESULTS

This checklist does not produce a score. It produces a governance profile. The following framework guides interpretation.

Result	Condition	Interpretation
Any ⚠ question answered No	Critical failure	The institution is not ready for deployment, regardless of all other responses. A single 'No' on a critical question identifies a governance gap that creates immediate liability or accountability exposure.
Any question answered Not Sure	Insufficient clarity	A 'Not sure' response on any question indicates that the governance condition in question has not been verified. Unverified conditions cannot be assumed to be met.
All questions answered Yes	Minimum conditions met	The institution meets the minimum governance conditions for deployment. This does not guarantee successful integration. It confirms that the foundational conditions are in place.

On the relationship between this checklist and deployment readiness

Completing this checklist with all 'Yes' responses does not guarantee successful AI integration. Integration, the process by which an institution genuinely changes how it makes decisions, requires sustained governance investment beyond the deployment date. This checklist identifies the floor, not the ceiling. What it confirms, when all conditions are met, is that the institution has addressed the governance gaps most consistently associated with institutional harm, legal exposure, and loss of public trust.

What this checklist does not assess

This checklist does not evaluate the technical capability of the AI system under consideration. It does not assess the accuracy of vendor claims. It does not determine whether the AI system is the right tool for the institution's objectives. It does not replace a full institutional readiness programme.

Its sole purpose is to determine whether the minimum governance conditions are in place. If they are not, no assessment of technical capability is sufficient to make deployment responsible.

The governing principle

The question is not whether the technology works.

The question is whether the institution is prepared to govern the consequences when it does not.

This checklist accompanies **AI Adoption in Institutions : What Must Be in Place Before You Deploy AI** *And Why Most Institutions Aren't Ready*

GUENIX

Developped by GUENIX DIGITAL
Author: Anastasie Guemtchuing Teuguia