

The Procurement File Is the Governance System

Heather Grizzle · August 30, 2026

This article argues that AI governance becomes real only in the procurement file, where institutions must document evidence justifying deployment decisions affecting disabled people.

1. Contents

1. Executive Summary
2. Governance Eventually Has to Reach the Transaction
3. The Procurement File as an Accountability Instrument
4. 1. Intended Use and Prohibited Use
5. 2. Affected-Population Analysis
6. 3. The Vendor Claim Record
7. 4. Independent Assurance
8. 5. Failure and Fallback Architecture
9. 6. Performance Boundaries and Known Unknowns
10. 7. Logging, Traceability, and Reconstruction
11. 8. Complaint, Contestability, and Redress
12. 9. Model Change and Revalidation
13. 10. Suspension and Exit Authority
14. From AI Policy to Institutional Evidence
15. NCG Policy Position
16. Closing Position

By Heather Grizzle, Novara Consulting Group

2. Executive Summary

Artificial intelligence governance is increasingly expressed through policies, committees, principles, inventories, and risk classifications. Those structures

matter. But when a public institution actually acquires and deploys an AI system, governance eventually has to become something much less abstract. It has to become a record.

Someone has to describe what the system is allowed to do. Someone has to decide what evidence is sufficient to support the vendor's claims. Someone has to identify the people who may be harmed when the system fails. Someone has to determine which failures require human intervention, which changes require reevaluation, and which conditions require the system to be suspended. If those decisions are not documented before deployment, an AI governance policy does not cure the omission. The procurement file is where institutional principles become operational commitments.

TL;DR

For AI systems affecting disabled people, the procurement file should contain enough evidence to answer a basic question after something goes wrong: why did the institution reasonably believe this system was appropriate for this use, with this population, under these conditions? If the answer cannot be reconstructed from the record, the governance system was incomplete before the system ever went live.

3. Governance Eventually Has to Reach the Transaction

A fallback is only a safeguard if the person who needs it can actually use it.

Heather Grizzle

Organizations are building increasingly sophisticated AI governance structures. They establish responsible AI principles, create review committees, maintain inventories, and classify systems according to risk. They adopt policies requiring fairness, transparency, accountability, privacy, accessibility, and human oversight. Then procurement begins. A vendor responds to a solicitation. Product demonstrations are scheduled. Accessibility documentation is collected. Security questionnaires are completed. Legal terms are negotiated. Business owners describe the efficiencies they expect. The technology clears the institution's established gates and enters service.

The existence of those gates can create considerable institutional confidence, but confidence is not the same thing as a defensible decision. The critical governance question is not merely whether the system passed through the required process. It is whether the resulting record contains the evidence necessary to justify the specific deployment that was approved.

That distinction becomes especially important when AI mediates access for disabled people. A system may summarize information, communicate with patients, evaluate applicants, translate language, generate instructional material, prioritize cases, assist employees, authenticate users, answer questions, or, increasingly, act on behalf of an institution. In each case, accessibility is not simply a characteristic of the interface. It is part of the decision architecture. A disabled person may encounter a system differently from the population on which it was primarily designed, tested, or validated, and an accessibility failure may alter the information that person receives, the choices available to them, their ability to respond, or the institution's interpretation of their behavior. The procurement record therefore

needs to establish more than whether accessibility was considered. It needs to establish what was considered, what was tested, what remains unknown, and what the institution decided to do about the uncertainty.

THE PROCUREMENT FILE AS AN ACCOUNTABILITY INSTRUMENT

4. The Procurement File as an Accountability Instrument

Procurement files are often understood primarily as administrative records. For AI, they should increasingly be understood as governance instruments. A mature AI procurement record should allow an independent reviewer to reconstruct the institution's reasoning without relying on the institutional memory of the people who participated in the purchase. That means the record should answer several questions:

1. What exactly was the system purchased to do?
2. What was it explicitly prohibited from doing?
3. Which populations were expected to encounter it?
4. What evidence supported the claims relied upon during approval, and who produced that evidence?
5. What limitations were known, and what limitations remained unknown?
6. What happened when the system was uncertain or wrong?
7. Who retained authority to override it?
8. How could an affected person challenge its output?
9. What would cause the institution to reconsider deployment?

10. Who had authority to stop it?

These are not questions for the incident investigation after a failure. They are procurement questions.

1. INTENDED USE AND PROHIBITED USE

5. 1. Intended Use and Prohibited Use

The first requirement is deceptively simple: define what the institution is buying. AI systems are frequently purchased through broad capability descriptions. A system may be described as an assistant, translator, accessibility tool, decision-support system, chatbot, agent, screening tool, or productivity platform. Those labels are inadequate governance specifications. The procurement record should identify the actual institutional function.

A communication tool used to help someone navigate a website presents a different risk from the same technology used during medical intake. An automated transcription system used to generate meeting notes presents a different risk from one used to create the official record of a disciplinary proceeding. A sign language recognition system used experimentally by an individual is not equivalent to one an institution presents as a substitute for another means of communication access. The same technology can therefore move between risk categories without any change to its underlying model, because the deployment context has changed.

DEFINE THE BOUNDARY BEFORE CONVENIENCE ERASES IT

The procurement file should state both the intended use and the prohibited use. Without the second, organizations create predictable scope expansion: a tool approved for a limited purpose becomes available, employees discover additional applications, and institutional use gradually exceeds the evidence on which the original approval rested.

2. AFFECTED-POPULATION ANALYSIS

6. 2. Affected-Population Analysis

Traditional technology procurement asks who will use the system. AI governance must also ask who will be affected by it, and those populations are not necessarily the same. An employee may operate an AI system while a patient, student, applicant, resident, customer, or benefits recipient experiences its consequences. Disabled people may be especially difficult to identify through conventional user analysis because disability frequently changes the pathway through a system rather than the nominal user category.

The procurement record should therefore identify foreseeable accessibility interactions. Does the system process speech? Does it generate or interpret language? Does it depend on vision, hearing, dexterity, cognition, timing, biometric characteristics, or a particular mode of communication? Does it infer intent or competence from behavior that disability may affect? Does an inaccessible output alter a person's ability to participate in a subsequent decision? These questions convert accessibility from a compliance artifact into a system-risk inquiry.

7. 3. The Vendor Claim Record

Every consequential vendor claim relied upon during procurement should have an evidentiary status. Not every claim requires independent testing, but institutions should distinguish among at least three categories. A capability is claimed when the vendor states that it exists. It is demonstrated when evidence shows that it has been observed under specified conditions. It is established for deployment when evidence reasonably supports reliance on it in the institution's intended context. Those categories are not interchangeable.

A demonstration can establish that a system is capable of producing an output. It does not necessarily establish how reliably it produces that output across populations, environments, or consequential use cases. Likewise, documentation can establish what a vendor represents about its system. It cannot independently establish that the representation is correct. The procurement file should preserve this distinction. Otherwise, assertions gradually become findings simply because they have been repeated in enough documents.

Established for deployment

A vendor claim status meaning evidence reasonably supports reliance on the capability in the institution's actual, intended context, not merely that the capability has been observed under demonstration conditions.

8. 4. Independent Assurance

Vendor evidence has a legitimate role in procurement. Vendors know their systems better than purchasers do and should be expected to disclose testing, architecture, limitations, known failure conditions, and relevant performance evidence. But the commercial relationship matters. The party seeking the contract should not be the only party determining whether the evidence is sufficient to award it.

Independent assurance does not mean recreating the vendor's engineering program. It means testing the claims that matter to the purchasing decision under conditions capable of disproving them. For accessibility-related AI, this may require evaluation involving people with the relevant disability and, where linguistic access is involved, qualified members of the language community. The objective is not ceremonial participation. It is evidentiary authority. The people capable of identifying a consequential failure must have a mechanism for causing that failure to affect the procurement decision. Otherwise, participation exists without governance.

5. FAILURE AND FALLBACK ARCHITECTURE

9. 5. Failure and Fallback Architecture

Every AI procurement should contain a written answer to a single question: what happens when the system does not work? "Human in the loop" is not an adequate answer. The record should identify the human, the trigger, the authority, the response time, and the alternative pathway available to the affected person.

For disabled users, the fallback itself must also be accessible. An AI communication system that fails and directs a Deaf person to make a telephone call has not

produced a fallback. An automated digital process that becomes inaccessible and requires a blind user to complete the same inaccessible interface has not transferred the person to human oversight in any meaningful sense.

6. PERFORMANCE BOUNDARIES AND KNOWN UNKNOWNNS

A fallback is only a safeguard if the person who needs it can actually use it.

Heather Grizzle

10. 6. Performance

Boundaries and Known Unknowns

Procurement documents tend to record what products can do. Governance records also need to state where confidence ends. AI performance is conditional. It may vary according to language, dialect, accent, signing style, lighting, camera position, disability, device, environment, domain terminology, interaction length, demographic characteristics, or countless other variables. Institutions should require vendors to disclose known performance boundaries and should separately document areas that have not been evaluated.

That distinction matters. Not evaluated is not the same finding as failed, but it is also not the same finding as safe. Uncertainty is not a procurement defect if it is recognized and governed. Undocumented uncertainty is.

7. LOGGING, TRACEABILITY, AND RECONSTRUCTION

11. 7. Logging, Traceability, and Reconstruction

When an AI-mediated interaction produces harm, the institution should be able to reconstruct what happened, and that requires decisions about logging before deployment. What system version was operating? What inputs were received? What output was generated? Was the output modified? Did a human review it? What information did that reviewer see? What action followed? Can the institution reproduce the event without retaining data that should never have been retained in the first place?

Traceability and privacy can create legitimate competing requirements. That conflict should be resolved through system design and governance rather than discovered during litigation, an accessibility complaint, or an internal investigation. The procurement file should document the resolution.

8. COMPLAINT, CONTESTABILITY, AND REDRESS

12. 8. Complaint, Contestability, and Redress

A person affected by AI needs more than a mechanism for reporting a technical problem. They may need to contest what the institution did because of the system, and that is a different function. If an AI system mistranslates communication, incorrectly characterizes behavior, produces inaccessible information, generates a false inference, or contributes to an adverse institutional action, the affected person needs a pathway to challenge the resulting decision.

The institution should know before deployment who receives the challenge, who investigates it, what evidence is preserved, whether the contested AI output remains operative during review, what remedies are available, and how quickly the

institution must respond.

A FEEDBACK FORM IS NOT A REDRESS MECHANISM

Redress requires an identified recipient, an investigation process, preserved evidence, a decision on whether the contested output stays live during review, defined remedies, and a response deadline.

9. MODEL CHANGE AND REVALIDATION

13. 9. Model Change and Revalidation

AI procurement creates a problem that traditional software contracts were not designed to handle particularly well: the product that was evaluated may not remain the product that is deployed. Models are updated. Guardrails change. Training data changes. Interfaces change. Third-party dependencies change. Vendors replace underlying models. Performance can improve in one domain and degrade in another without the purchasing institution initiating any change itself.

The procurement file therefore needs a change-control rule. The relevant question is not whether the vendor is permitted to update its product. It is which changes invalidate the evidence on which the institution relied. Contracts should identify material changes requiring notice, reevaluation, renewed acceptance testing, or approval before continued use in specified contexts. Otherwise, an institution can perform rigorous assurance at acquisition and still find itself governing a materially different system six months later using evidence produced for the previous one.

14. 10. Suspension and Exit Authority

Governance without stop authority is advisory. Someone within the institution must have the authority to restrict, suspend, or terminate an AI deployment when evidence no longer supports its use, and the triggering conditions should not have to be invented during a crisis.

- ☐ Repeated accessibility failures
- ☐ Undisclosed material model changes
- ☐ Loss of required human oversight
- ☐ Significant divergence from validated performance
- ☐ Unresolved complaints
- ☐ New evidence of population-specific harm
- ☐ Failure to provide contractually required information
- ☐ Expansion beyond the approved use case

The procurement record should identify both the triggers and the decision authority.

It should also address exit. Who owns or receives institutional data when the relationship ends? What records must be retained? What must be deleted? How does the institution restore the previous service pathway? What happens to people who have become dependent on the AI-mediated process? Exit planning is not merely vendor management. It is continuity-of-access planning.

15. From AI Policy to Institutional Evidence

The emerging challenge for AI governance is not a shortage of principles. It is translation. Fairness has to become an evaluation criterion. Transparency has to become a disclosure requirement. Human oversight has to become an assigned authority. Accessibility has to become a deployment condition. Accountability has to become a named office and a documented remedy. Monitoring has to become a trigger for action. And responsible AI eventually has to become something a procurement officer can place in a file.

This is particularly important in public institutions because procurement is where abstract institutional commitments encounter money, contracts, vendors, statutory duties, and people who cannot simply choose another government, school, hospital, court, or public service when the technology does not work for them. The procurement file therefore performs a function much larger than recordkeeping. It establishes the institution's theory of acceptable risk.

NCG POLICY POSITION

16. NCG Policy Position

Novara Consulting Group holds that AI systems affecting disabled people should not be approved solely because they have satisfied an organization's general AI governance process. The procurement record should contain sufficient evidence to reconstruct and defend the specific deployment decision.

AT MINIMUM, THE RECORD SHOULD ESTABLISH

- The intended and prohibited uses of the system
- The populations foreseeably affected
- The evidentiary status of consequential vendor claims
- The basis for independent assurance where warranted
- Known failure conditions and untested conditions
- Accessible fallback and human-review pathways
- Logging and traceability requirements
- Complaint, contestability, and redress procedures
- Material-change and revalidation triggers
- Institutional authority to restrict, suspend, or terminate deployment

The standard is not perfect knowledge. No procurement process can eliminate uncertainty. The standard is governed uncertainty: knowing what has been established, what has not, what evidence justified proceeding anyway, and who bears responsibility for the decision.

CLOSING POSITION

17. Closing Position

After an AI failure, institutions often ask whether the vendor should have warned them. Sometimes the answer will be yes. But public accountability requires another question: what evidence did the institution require before it agreed to proceed?

That question cannot be outsourced. The vendor can supply evidence. An assessor can evaluate it. A committee can advise. Disabled stakeholders can identify failures that others cannot see. Lawyers can negotiate the allocation of contractual risk. The institution still makes the decision, and its procurement record should prove that it knew what decision it was making.

An AI system is not governed because an organization has an AI policy. It is governed when the institution can show why this system, for this purpose, affecting these people, under these conditions, was permitted to operate.

Heather Grizzle

The procurement file is where that accountability begins.