

Who Audits the Audit Trail? AI-Generated Conformance Claims and the Procurement Problem

Heather Grizzle · July 15, 2026

Generative AI has changed accessibility procurement. Learn why VPATs and Accessibility Conformance Reports must be supported by verifiable evidence.

The history of public procurement can be read, in part, as the history of documentary trust. Governments rarely purchase software after examining its source code. Universities do not independently inspect every security control claimed by a learning management system, nor do hospitals routinely recreate the accessibility testing performed by a medical software vendor. Procurement depends instead upon documentary representations that stand between the purchaser and the product itself. Certifications, technical reports, audit findings, conformity assessments, and Accessibility Conformance Reports exist because direct verification is expensive, specialised, and frequently impractical. Institutions therefore purchase not only products but also the evidence offered in support of those products.

That arrangement has always required judgment. A VPAT was never intended to constitute proof that accessibility existed. It was designed as a structured representation of a vendor's evaluation against recognised accessibility criteria. Procurement officials understood that distinction even when they relied heavily upon the document itself. The report carried weight not because it possessed independent legal authority but because preparing it ordinarily required someone to spend considerable time evaluating the product. The effort invested in producing the document became, however imperfectly, a proxy for the effort invested in producing the underlying evidence.

Generative artificial intelligence disrupts that relationship without altering the document's appearance.

An organization can now produce accessibility documentation that is technically fluent, internally consistent, professionally formatted, and persuasive in remarkably little time. Standards are cited correctly. Technical terminology is used appropriately. The structure resembles thousands of conformance reports that procurement officials have reviewed over the past two decades. None of those characteristics, however, reveal whether the document reflects extensive accessibility testing, limited internal review, or no meaningful evaluation at all. The sophistication of the writing has become increasingly detached from the sophistication of the evidence.

This is not an argument against artificial intelligence. Documentation has always been an act of translation. Engineers translate software behaviour into technical findings. Accessibility specialists translate those findings into procurement language. Procurement officials translate technical language into purchasing decisions. Generative AI simply accelerates one stage of that translation process. The difficulty emerges when translation is mistaken for verification. A well-written report may accurately describe a rigorous evaluation, but it may also describe assumptions that have never been independently examined. The document itself provides no reliable means of distinguishing between those possibilities.

The sophistication of the writing has become increasingly detached from the sophistication of the evidence.

The consequence is more significant than it first appears. Procurement files now contain documents that often look substantially alike despite representing fundamentally different evidentiary conditions. One Accessibility Conformance Report may summarize independent testing conducted by experienced accessibility professionals using both automated and manual methodologies. Another may represent an internal engineering assessment that

has never been externally reviewed. A third may accurately document extensive testing while relying upon generative AI to improve organization and readability. A fourth may have been generated almost entirely from product specifications and marketing materials without the underlying software ever being evaluated against accessibility standards. Once reduced to a professionally formatted report, each occupies the same evidentiary space within the procurement record despite resting upon very different foundations.

The traditional procurement question has therefore become less useful than it once was. For many years, accessibility evaluations centred upon whether a product conformed to technical standards. Increasingly, however, the more consequential inquiry concerns the basis upon which that conclusion rests. The existence of a conformance claim is no longer particularly informative when the cost of producing persuasive documentation has fallen so dramatically. Procurement officials instead require information about the origin of the claim itself: who performed the evaluation, which methodologies were employed, whether manual testing accompanied automated scanning, whether assistive technology users participated, whether findings can be independently reproduced, and whether the supporting evidence remains available for external review. These questions concern evidence rather than accessibility, yet they increasingly determine whether accessibility claims deserve institutional confidence.

QUESTIONS ABOUT THE CLAIM, NOT THE PRODUCT

Who performed the evaluation, which methodologies were employed, whether manual testing accompanied automated scanning, whether assistive technology users participated, whether findings can be independently reproduced, and whether the supporting evidence remains available for external review.

Recent procurement guidance has begun to reflect this broader shift. Rather than treating documentary submission as the endpoint of due diligence, public-sector guidance increasingly emphasizes substantiation, reproducibility, methodological transparency, and independent verification. The significance of an Accessibility Conformance Report therefore lies less in its existence than in its ability to direct procurement officials toward evidence capable of surviving meaningful scrutiny. The report becomes valuable not because it asserts compliance, but because it explains why the assertion should be believed.

The implications extend well beyond accessibility procurement. Artificial intelligence has substantially reduced the cost of producing persuasive institutional documents across nearly every domain of governance. Security assessments, risk analyses, governance statements, implementation plans, privacy documentation, and procurement narratives can all be drafted with increasing fluency and decreasing effort. Yet the practical work upon which those documents depend remains resistant to similar acceleration. Independent testing still requires qualified evaluators. Validation still depends upon reproducible methodology. Evidence still emerges from observation rather than composition. The imbalance created by these developments is therefore structural rather than technological. The production of representations has accelerated far more rapidly than the production of the facts those representations describe.

It is within this context that evidentiary governance becomes increasingly important. The Sign Language Access Trust (SLAT) Framework was developed from the premise that documentation should not be evaluated independently of the institutional processes that produce it. Rather than treating vendor assertions as self-authenticating, the framework examines governance structures, validation methodology, transparency practices, disclosure quality, and the provenance of supporting evidence. The central question is not whether a claim has been made but whether sufficient evidence exists for a reasonable purchaser to rely upon that claim when making consequential decisions affecting accessibility, public expenditure, and legal accountability.

Public procurement has long depended upon documentary artefacts because institutions cannot independently verify every technical assertion presented to them. That dependence is unlikely to disappear. What has changed is the relationship between documentation and the effort historically required to produce it. As artificial intelligence continues to lower the cost of persuasive technical writing, procurement practice will necessarily place greater emphasis upon the quality, traceability, and independence of the evidence underlying those documents. The future credibility of accessibility procurement will depend less upon the sophistication of conformance reports than upon the ability of institutions to distinguish documentary representation from demonstrable fact.

Author's Note

Heather M. Grizzle-Odland is the Founder and Chief Executive Officer of Novara Consulting Group. Her work focuses on AI accessibility governance, sign language AI procurement risk, evidentiary standards, and Deaf-led accountability. She is also the developer of the Sign Language Access Trust (SLAT) Framework, a structured evaluation model for assessing the governance, validation, transparency, and procurement readiness of sign language AI systems.